

THÈSE DE DOCTORAT

de

L'UNIVERSITÉ PARIS-SACLAY

École doctorale de mathématiques Hadamard (EDMH, ED 574)

Établissement d'inscription : Université Paris-Sud

Laboratoire d'accueil : Laboratoire de mathématiques d'Orsay, UMR 8628 CNRS

Spécialité de doctorat : Mathématiques appliquées

Claire BRÉCHETEAU

Vers une vision robuste de l'inférence géométrique

Date de soutenance : 24 septembre 2018

Après avis des rapporteurs :

GÉRARD BIAU (Sorbonne Université)

KATHRYN HESS BELLWALD (École Polytechnique Fédérale de Lausanne)

Jury de soutenance :

GÉRARD BIAU (Sorbonne Université) Rapporteur

FRÉDÉRIC CHAZAL (Inria Saclay) Codirecteur de thèse

KATHRYN HESS BELLWALD (École Polytechnique Fédérale de Lausanne) Rapporteur

MARC HOFFMANN (Université Paris-Dauphine) Président du jury

PASCAL MASSART (Université Paris-Saclay) Directeur de thèse

QUENTIN MÉRIGOT (Université Paris-Saclay) Examineur

Remerciements

Cette thèse découle d'une nouvelle collaboration entre Pascal Massart et Frédéric Chazal, des personnalités qui ont toujours su s'accorder. À mes deux directeurs de thèse, je souhaite vous exprimer mon entière gratitude et ma reconnaissance pour avoir su m'octroyer votre confiance, et ce malgré les fluctuations truculentes qui vous ont maintenus en haleine jusqu'à la fin. Merci de m'avoir encouragée à suivre mes idées jusqu'au bout, quitte à m'éloigner de la voie que m'aviez esquissée. Merci aussi d'avoir su être présents aux moments charnières. Fred, tu n'as pas rechigné à lire ces "romans 'presque' littéraires"... merci pour tes petits conseils bien avisés. Ta petite recette magique, simple mais efficace pour la rédaction, je m'efforcerai de la suivre désormais. Pascal, je me souviendrai toujours de ces minutes passées dans ton bureau à essayer d'exprimer des idées avec plus ou moins de clarté. Tu ne t'es jamais esquivé lorsqu'il s'agissait d'écouter et essayer de saisir les sens cachés. Merci à vous deux. Je garderai un excellent souvenir de nos échanges.

Un grand merci à mes deux rapporteurs, Kathryn Hess et Gérard Biau, pour avoir accepté de rapporter ma thèse, en plein milieu d'un été particulièrement ensoleillé. Merci également à Marc Hoffmann et à Quentin Mérigot pour avoir accepté de faire partie de mon jury de thèse.

Je tenais également à remercier l'école doctorale et surtout Stéphane Nonnenmacher et Frédéric Paulin pour leur investissement et pour avoir su rendre les tâches administratives liées à la thèse beaucoup moins stressantes. Merci également aux assistantes pour leur travail précieux tout au long de ces trois années.

Cette thèse aurait été largement amputée quant à son contenu si tu n'avais pas été là, Clément. Un grand merci pour tous ces échanges fructueux. Tu m'as beaucoup appris. Ça a été un réel plaisir d'échanger avec toi et d'aller au bout d'idées parfois bien farfelues. Aurélie, mes pensées vont également vers toi. Je pense en particulier à la bataille acharnée que nous avons menée ensemble au milieu de l'été pour cette quête de données, entre deux plongeurs.

Remontons un peu le temps. Luck, par tes conseils, tu m'as guidée vers un point où mon imagination me voyait une dizaine d'années auparavant. Je t'en remercie.

Merci aux professeurs qui m'ont construite sur le plan des mathématiques. Je pense en particulier à François Sauvageot et à Bernard Luron qui ont été présents dans mes premiers vrais contacts avec les mathématiques. M. Sauvageot, vous aimez jouer avec les mathématiques

et les rendre appréciables aux yeux des autres avec une chaleur humaine dont la réputation n'est plus à faire. M. Luron, je crois me souvenir que "[vous vous souviendrez] toujours de votre avant-dernière copie", moi je me souviendrai de votre bienveillance, même envers les moins vifs.

L'arrivée à la faculté d'Orsay a été ensuite une vraie révélation. Je retiendrai ainsi de mes années à l'Université Paris-Sud cet enseignement d'excellence, et ses enseignants remarquables. Vous avez su me laisser entrevoir la beauté des objets mathématiques, tout en préservant une rigueur qui m'est plaisante. En particulier, je souhaite adresser mon respect et ma gratitude à Dominique Hulin, vous nous avez tellement donné lors de notre arrivée en L3 et j'ai adoré votre cours de topologie générale, mais aussi à Jean-François Le Gall et à Nicolas Curien pour nous avoir laissé percevoir un monde plein de beauté avec élégance.

Lors de mes années de thèse, j'ai été à cheval entre deux laboratoires, le LMO à Orsay (ainsi que l'équipe Inria Select dont je salue les membres avec une pensée particulière pour Élisabeth Gassiat, Sylvain Arlot, Christophe Giraud et Gilles Celeux...), et l'équipe Inria Datashape sur le plateau de Saclay (qui a succédé à l'équipe Géométrica). Une pensée chaleureuse me vient pour Jeanne, Martin, Benjamin, Luc et Thomas, les Guillaume, Salim et tous les autres doctorants d'Orsay qui ont dû me supporter dans la dernière ligne droite. Je garderai un bon souvenir de ces quelques journées où nous finissions entassés à 4 dans des bureaux de 3 pour s'apporter un soutien mutuel, laissant vides de grands bureaux neufs.

Merci aux collègues de Datashape, je pense à Steve, Eddie, Raphaël, Jérémy, Tullia, Harry, Paweł, Thomas, Mathieu, Vincent, Bertrand, Théo, Miro... et les nouveaux qui ont pris la relève. Sans oublier Ilaria et Nicolas avec qui j'ai eu le plaisir de partager un bureau.

En particulier, je tenais à te remercier Marc ; tu étais toujours partant pour que je te raconte mes petites avancées et aïde de mes figures. Je garderai toujours un excellent souvenir de nos échanges. Merci aussi pour tout le temps que tu as consacré à me donner des conseils pour mes codes en C++. En passant, je salue Vincent pour son cours d'initiation à ce langage, fort bien mené, et pour son calme et sa gentillesse indéfectibles !

Parmi les membres qui ont vécu à Datashape, je souhaitais aussi et surtout saluer Bertrand. Merci Bertrand de m'avoir fait confiance, j'ai hâte de poursuivre ce postdoc avec toi pour d'autres aventures statistiques.

Je salue également les membres de la partie niçoise de l'équipe Datashape, menés par Jean-Daniel, avec qui nous avons passé quelques semaines bien enrichissantes.

Cette thèse a été agrémentée d'un court séjour à l'étranger. Many thanks to Gosia for manifesting interest to this work, and allowing to share together for 3 weeks in Wrocław. You helped me a lot, and I'll always remind of your kindness and welcoming side. By the way, I'd like to greet cheerfully my friends from Wrocław, and particularly Rafał. Many thanks for allowing me to discover from the inside, a new culture that I love.

Écrire une thèse, surtout en dernière année, requiert un environnement stable. En ce sens, je tenais à saluer mes proprios pour n'avoir pas rechigné à loger une doctorante d'âge avancé.

Et parce que l'esprit et le corps ont du mal à fonctionner l'un sans l'autre, je tenais à remercier plus que chaleureusement le capitaine Rachid. Merci de m'avoir intégrée dans l'équipe de foot d'Orsay. Vous m'avez tous permise de m'évader et d'oublier un peu le stress. Nos matchs passionnés, les fives et les soirées vont me manquer. Au sanglier des Ardennes, j'ai apprécié partager ces moments dans ce cadre plus sportif qu'est celui du ballon rond. Un petit salut au passage aux anciens de l'équipe (de foot) du CMAP, célèbre pour sa combativité lors des célèbres tournois annuels inter-labos.

Un esprit sain dans un corps sain, c'est aussi toi qui me l'a appris, Rémi. Merci pour les infusions miracles, les voyages et tous ces échanges philosophiques et plus. Tu m'as redonné la confiance et l'énergie nécessaire pour achever cet ouvrage. Je ne l'oublierai pas, je te dois énormément.

La vie autour de cette thèse aurait manqué de saveur sans la présence de mes autres amis que je salue avec ferveur et que je remercie. Antoine, pour m'avoir initiée aux douceurs de la vie parisienne.

Ludo, Lia, Rossica, Benjamin, Pierre et Lilian... pour tous les bons moments mathématiques et non mathématiques qu'on a pu partager ensemble. Et un merci tout particulier et plein d'affection à Benjaminou, le roi de la wifi et des kouign-amis !

Cormac, pour m'avoir ouverte aux variétés finslériennes et ne pas avoir été fermé à l'idée que des filles puissent jouer au football en ayant accepté de m'intégrer dans l'Équipe de foot du CMAP. Merci pour ton amitié sans failles. Et accessoirement, merci d'avoir relu une bonne partie de l'anglais de mes écrits, même si la compensation était digne d'intérêt.

Alain, parce que cette thèse représente l'aboutissement d'un chemin que nous avons parcouru longtemps ensemble, laisse moi t'exprimer mon affection et ma gratitude. Je pense aussi à tous ces élèves et étudiants qui auront la chance d'accéder à l'enseignement d'un professeur si bienveillant, impliqué, calme et vraiment doué, disons le.

Enfin et surtout, si "au no 16 de la rue Grande-Pohulanka, à Wilno, habitait un certain M. Piekielny", j'aimerais profiter de ces remerciements de thèse pour annoncer que quelque part dans les Mauges, habitent une poignée de gens simples et généreux qui participent à faire tourner le monde, dans l'ombre. Grâce à eux et à tous les autres eux, je fais partie de ceux qui ont eu la possibilité de vivre quelques années de leur passion. En particulier, j'aimerais remercier plus que chaleureusement ma famille. Vous avez toujours su me soutenir.

Contents

1	Introduction	2
1.1	Notions fondamentales et État de l’art	4
1.1.1	Une introduction à la comparaison de mesures	4
1.1.2	Une introduction aux méthodes de partitionnement et de quantification	9
1.1.3	Une courte introduction à l’inférence géométrique et topologique	15
1.2	Contributions et Contenu	18
1.3	Présentation détaillée des chapitres	23
2	Similarity of distributions according to the Gromov-Wasserstein distance	28
2.1	Discriminative tools for comparing probability distributions	29
2.1.1	From comparison of compact sets to comparison of probability distributions	29
2.1.2	Equality of probability distributions up to rigid transformations	35
2.2	Statistical test of similarity in the Gromov-Wasserstein sense	40
2.2.1	A test of asymptotic level α	42
2.2.2	The power of the test	45
2.2.3	The algorithm	46
2.2.4	A discussion of other methods	49
2.2.5	On the selection of the parameters	52
2.3	Proofs	53
2.3.1	Proofs for Section 2.1	53
2.3.2	Proofs for Section 2.2	61
	Appendices	73
2.A	A characterization of metric-measure spaces up to an isomorphism – On the proof of a theorem from Gromov	73
2.A.1	The Theorem of identifiability of measures up to an isomorphism	73
2.A.2	Proofs	75
3	Trimmed criteria for clustering with applications to mixture models	78
3.1	Trimmed clustering with Bregman divergences	79
3.1.1	Context and motivation	80

3.1.2	Contribution	81
3.1.3	A framework for trimmed clustering with Bregman divergences	82
3.1.4	Algorithm	87
3.1.5	Experiments	88
3.2	Heteroscedastic Gaussian mixture clustering with the trimmed log-likelihood	94
3.2.1	The trimmed log-likelihood - a notion related to the distance to measure	95
3.2.2	A new clustering method based on the trimmed log-likelihood and a modified version of the Lloyd's algorithm	99
3.2.3	Experiments	105
3.3	Proofs	109
3.3.1	Proofs for Section 3.1	109
3.3.2	Proofs for Section 3.2	120
Appendices		126
3.A	Maximal deviation bounds for empirical processes – Background and first results	126
3.A.1	From bounded difference inequality to Rademacher averages	127
3.A.2	\mathcal{F} , a family of indicators	129
3.A.3	\mathcal{F} , a family of bounded functions	130
4	A new quantization method based on the distance to measure	142
4.1	The deterministic framework	145
4.1.1	The k -PDTM: a coresset for the DTM	145
4.1.2	Proximity to the DTM	148
4.1.3	Wasserstein stability for the k -PDTM	149
4.1.4	Geometric inference with the k -PDTM	149
4.2	The statistical framework	150
4.2.1	An algorithm for the empirical k -PDTM	151
4.2.2	Proximity between the k -PDTM and its empirical version	152
4.2.3	Numerical illustrations	153
4.3	The set of local means – A strong connection with the distance-to-measure	159
4.3.1	Local means – From balls to half-spaces	160
4.3.2	Convexity of the set of local means	162
4.3.3	Duality, set of local means and distance-to-measure	163
4.4	Proofs	167
4.4.1	Proofs for Section 4.1	167
4.4.2	Proofs for Section 4.2	172
4.4.3	Proofs for Section 4.3	181
Bibliography		188
Index		197

Introduction

Le volume de données disponibles est en perpétuelle expansion. Il est primordial de fournir des méthodes efficaces et robustes permettant d'en extraire des informations pertinentes. Le but étant d'en réduire le volume afin de pouvoir les traiter plus efficacement ou de pouvoir comparer plus facilement deux ensembles de données. Souvent ces données peuvent être représentées sous la forme de nuages de points dans un certain espace muni d'une métrique, e.g. l'espace Euclidien \mathbb{R}^d . Cette représentation rend l'étude mathématique de cette source d'information plus aisée. En particulier, les théories des statistiques, des probabilités et de la mesure fournissent des bases mathématiques solides à l'étude de ce type de données. À ces disciplines se sont ajoutées plus récemment la géométrie et la topologie, qui ont déjà fait leurs preuves en analyse de données, on pense en particulier au domaine naissant de l'analyse topologique des données (TDA).

Parmi les questions naturelles que l'on peut se poser lorsque l'on a accès à des données, trois d'entre elles sont abordées dans cette thèse. La première concerne la *comparaison de deux ensembles de points*. Comment décider si deux nuages de points sont issus de formes ou de distributions similaires ? Il convient de se questionner sur la notion de "similarité" de deux ensembles ou deux mesures. L'approche proposée repose sur l'étude de descripteurs de données, stables et robustes à certains types de bruit mais également discriminants. L'un des objectifs est de construire un test statistique permettant de décider si deux nuages de points sont issus de distributions égales (modulo un certain type de transformations e.g. symétries, translations, rotations...). À notre connaissance, aucun test statistique de ce type n'a été étudié dans la littérature.

La seconde question concerne la *décomposition d'un ensemble de points en plusieurs groupes*. Étant donné un nuage de points, comment faire des groupes pertinents ? Souvent, cela consiste à choisir un système de k représentants et à associer chaque point au représentant qui lui est le plus proche, en un sens à définir. Il s'agit du problème de partitionnement des données. Dans cette thèse, les méthodes que nous introduisons et étudions sont adaptées à des données échantillonnées selon un mélange¹ de lois, avec éventuellement quelques données aberrantes. Dans un premier temps, nous traitons le cas des mélanges de lois de familles exponentielles, pour lesquelles la densité peut s'écrire de façon naturelle en fonction d'une divergence de

¹On appelle *mélange de loi* tout distribution P de type $P = \sum_{i=1}^k \alpha_i P_i$. En particulier, P est un mélange de lois normales lorsque les P_i sont des lois normales.

Bregman. C'est le cas des lois Gaussiennes de matrice de covariance fixée, des lois de Poisson, binomiales, Gamma... Dans un second temps, nous traitons le cas des mélanges Gaussiens hétéroscédastiques². La particularité des méthodes proposées est qu'elles dépendent d'un paramètre de trimage h , à sélectionner à partir des données. Seule une proportion $1 - h$ des données sera partitionnée. L'étude des méthodes proposées repose sur des outils géométriques tels que la distance à la mesure³ ou les réseaux de Voronoï. En particulier, elle repose sur des versions continues du critère des k -moyennes, du critère de quantification avec une divergence de Bregman et de la log-vraisemblance trimagee ; ces versions continues peuvent toutes être vues comme des généralisations de la fonction distance à la mesure.

Enfin, lorsque les données n'ont pas naturellement une structure en k groupes, par exemple, lorsqu'elles sont échantillonnées à proximité d'une sous-variété de \mathbb{R}^d , une question plus pertinente est de construire un système de k représentants, avec k grand, à partir duquel on puisse retrouver la sous-variété. Cette troisième question recouvre le problème de la *quantification* d'une part, et le problème de l'*approximation de la distance à un ensemble* d'autre part. De nouveau, les méthodes proposées sont adaptées à des données éventuellement corrompues par quelques données aberrantes, et reposent sur des outils géométriques. Plus précisément, la construction d'un tel système de représentants revient à fournir une bonne approximation de la fonction distance à la mesure.

Les réponses que nous apportons à ces trois questions dans cette thèse sont de deux types, théoriques et algorithmiques. En particulier, les méthodes proposées reposent sur des objets continus, définis directement à partir de mesures de probabilité et de sous-mesures. Des études statistiques permettent de mesurer la proximité entre les objets empiriques et les objets continus correspondants. Ces méthodes sont faciles à implémenter en pratique lorsque des nuages de points sont à disposition. L'outil principal utilisé dans cette thèse est la fonction distance à la mesure. À l'origine, elle a été introduite pour stabiliser les méthodes d'analyse topologique des données. C'est-à-dire, pour que les informations topologiques contenues dans des échantillons de points ne soient pas perturbées par la présence de quelques données aberrantes. Ses intérêts sont multiples. Elle permet d'inférer le support d'une distribution à partir de données éventuellement bruitées. Son écriture comme fonction puissance⁴ permet d'établir un lien étroit avec le critère des k -moyennes. Aussi, ses propriétés de stabilité par rapport aux métriques de Wasserstein en font un outil particulièrement adapté à l'échantillonnage et facilite l'étude statistique des objets d'intérêt.

Les apports de cette thèse recourent plusieurs domaines. Celui de la comparaison de distributions, le partitionnement des données et l'inférence géométrique. Nous commençons par introduire les notions fondamentales et l'état de l'art relatifs à ces domaines.

²On parle de mélange hétéroscédastique lorsque les variances des distributions du mélange diffèrent.

³La distance à la mesure est une généralisation de la fonction distance à un compact, aux mesures de probabilité. Elle dépend d'un paramètre de trimage h . Pour une distribution uniforme sur un ensemble de n points \mathbb{X}_n dans \mathbb{R}^d , la distance à la mesure coïncide avec le critère des moindres carrés trimageés. Elle est alors définie en tout point de \mathbb{R}^d comme la moyenne des carrés des distances à ses $q = nh$ plus proches voisins dans \mathbb{X}_n .

⁴Une fonction puissance est une fonction définie sur \mathbb{R}^d , de la forme $x \mapsto \inf_{i \in I} \|x - c_i\|^2 + \omega_i$, avec I un ensemble quelconque, $c_i \in \mathbb{R}^d$ et $\omega_i \in \mathbb{R}$.

1.1 Notions fondamentales et État de l'art

1.1.1 Une introduction à la comparaison de mesures

Des métriques ou pseudo-métriques pour comparer des mesures

Borel et Jordan sont à l'origine de la théorie de la mesure des ensembles ([Bor98] en 1898) et de la théorie moderne du calcul des probabilités dans le cadre des σ -algèbres⁵ ou tribus ([Bor30] en 1930).

Une fois les σ -algèbres définies, la théorie moderne des probabilités repose sur l'axiomatique de Kolmogorov introduite par le mathématicien russe en 1933 dans [Kol33]⁶. Un *espace probabilisé* est un triplet (Ω, \mathcal{F}, P) constitué d'un ensemble Ω muni d'une σ -algèbre \mathcal{F} et d'une mesure de probabilité P . Une *mesure de probabilité* (*distribution* ou *loi*) est une application définie sur \mathcal{F} satisfaisant les trois axiomes suivants :

- $\forall A \in \mathcal{F}, P(A) \in [0, 1]$,
- $P(\Omega) = 1$,
- P est σ -additive au sens où si les A_n sont deux-à-deux disjoints, $P(\bigcup_{n \in \mathbb{N}} A_n) = \sum_{n=1}^{+\infty} P(A_n)$.

En particulier, $P(\emptyset) = 0$.

On dira qu'une *variable aléatoire*⁷ X suit la loi P (noté $X \sim P$) si la probabilité que X soit à valeurs dans A vaut $P(A)$. Pour f une fonction mesurable⁸, on notera $Pf(u)$ ou $\mathbb{E}_{X \sim P}[f(X)]$ l'intégrale de f par rapport à la distribution P .

Plus généralement, on dit que P est une *mesure positive finie* si P est à valeur dans \mathbb{R}^+ et P est σ -additive. En particulier, $P(\Omega) < \infty$. Un exemple de mesures positives qui ne sont pas nécessairement des mesures de probabilité sont les *sous-mesures*⁹ d'une distribution P .

Une première idée naturelle pour comparer deux distributions consiste à calculer la différence maximale entre leurs valeurs prises sur les éléments de la tribu. La métrique correspondante est la distance en variation totale d_{VT} définie pour toutes mesures de probabilité P et Q sur le même espace mesurable (Ω, \mathcal{F}) par :

$$d_{VT}(P, Q) = \sup_{A \in \mathcal{F}} |P(A) - Q(A)|.$$

Le théorème de Pinsker relie la distance en variation totale à la divergence de *Kullback-Leibler* ou *entropie relative*. Il s'agit d'une autre mesure de dissimilarité entre deux distributions P et Q fréquemment utilisée en statistiques.

Lorsque P est une distribution sur (Ω, \mathcal{F}) et Q sur (Ω', \mathcal{F}') , on dit qu'une mesure de probabilité π définie sur $\Omega \times \Omega'$ ¹⁰ est un *plan de transport* ou *couplage* entre P et Q si $\pi(A \times \Omega') = P(A)$ et $\pi(\Omega \times A) = Q(A)$ pour tout Borélien A . Ainsi, un plan de transport π est la loi d'un couple de variables aléatoires (X, Y) tel que $X \sim P$ et $Y \sim Q$. On note $\Pi(P, Q)$ l'ensemble des plans de transport entre P et Q .

⁵Une σ -algèbre ou tribu est définie à partir d'un ensemble Ω comme une collection \mathcal{F} de parties de Ω qui vérifie les trois axiomes suivants : $\Omega \in \mathcal{F}$ et $\emptyset \in \mathcal{F}$; $\forall A \in \mathcal{F}, A^c \in \mathcal{F}$, où $A^c = \Omega \setminus A$ est le complémentaire de A dans Ω ; et $\forall n \in \mathbb{N}, \forall A_n \in \mathcal{F}, \bigcup_{n \in \mathbb{N}} A_n \in \mathcal{F}$.

⁶[Kol33] est un ouvrage en allemand intitulé : Notions fondamentales du calcul des probabilités

⁷Une variable aléatoire $X : (\Omega_0, \mathcal{F}_0, P_0) \rightarrow (\Omega, \mathcal{F}, P)$ est une *fonction mesurable* (au sens où $X^{-1}(A) \in \mathcal{F}_0$ pour tout $A \in \mathcal{F}$) qui satisfait pour tout $A \in \mathcal{F}, P_0(X^{-1}(A)) = P(A)$.

⁸Voir la note précédente.

⁹Une sous-mesure \tilde{P} d'une mesure de probabilité P est mesure positive \tilde{P} telle que : $\forall A \in \mathcal{F}, \tilde{P}(A) \leq P(A)$.

¹⁰ $\Omega \times \Omega'$ est muni de la tribu produit $\mathcal{F} \otimes \mathcal{F}'$

La distance en variation totale se réécrit en termes de plans de transports comme suit :

$$d_{\text{VT}}(P, Q) = \inf_{\pi \in \Pi(P, Q)} \mathbb{E}_{(X, Y) \sim \pi} [\mathbb{1}_{X \neq Y}]. \quad (1.1)$$

La distance en variation totale est une *métrique de transport* au sens où elle consiste à minimiser l'espérance d'une fonction de coût $(x, y) \mapsto \mathbb{1}_{x \neq y}$ par rapport à un plan de transport. Cette métrique n'est pas adaptée à l'échantillonnage ni à la convergence en loi lorsque P est une distribution à densité par rapport à la mesure de Lebesgue dans \mathbb{R}^d . Lorsque P_n désigne la distribution empirique¹¹ associée à P , alors $d_{\text{VT}}(P, P_n)$ vaut toujours 1. Pourtant, dans ce cas, le théorème de Varadarajan [Var58b] assure que P_n converge en loi vers P presque sûrement.

Désormais, on travaille avec $\Omega = \mathcal{X}$, un espace¹² muni d'une métrique δ et de la tribu Borélienne¹³ $\sigma(\mathcal{X})$. Lorsque cette métrique le rend séparable et complet, on dit que \mathcal{X} est un *espace Polonais*. Concernant les propriétés intéressantes de théorie de la mesure et des espaces Polonais, on pourra se référer à l'ouvrage de Cohn [Coh80].

On peut remplacer la fonction de coût $(x, y) \mapsto \mathbb{1}_{x \neq y}$ par $(x, y) \mapsto \delta(x, y)^p$ pour un $p \geq 1$ dans (1.1). On retrouve alors les métriques de Wasserstein, aussi connues sous le nom de Mallows, Kantorovich-Monge-Rubinstein ou distance du cantonnier (earth mover distance en anglais). Elles ont été introduites en 1969 par le mathématicien russe Vasseršteïn et leur appellation a été donnée par Dobrushin en 1970 ([Dob70]), voir aussi [Vil09]. Elles sont donc définies par :

$$W_p^p(P, Q) = \inf_{\pi \in \Pi(P, Q)} \mathbb{E}_{(X, Y) \sim \pi} [\delta^p(X, Y)]. \quad (1.2)$$

Comme la distance de Lévy-Prokhorov définie en 1956 par Prokhorov [Pro56] à partir des travaux de Lévy en 1937 [Lév37], les métriques de transport W_p ainsi obtenues sont adaptées à la convergence en loi (ou convergence faible). De plus, elles sont également adaptées à la convergence des moments d'ordre p associés, $P\delta^p(u, x_0)$ pour $x_0 \in \Omega$, voir [BL14, Théorème 2.6], [Vil03, p.212], ou [RAW09].

Parmi les pseudo-distances permettant de mesurer la disparité entre deux distributions, on peut également citer les métriques de probabilités intégrées ou encore mesures de discrédance moyenne maximale (voir [eE53, Mü97, GMBJR⁺12]). Des cas particuliers de telles mesures de dissimilarité sont la métrique de Kolmogorov-Smirnov et les distances de Wasserstein ou de Kantorovich et Rubinstein (voir [KR58]).

En inférence géométrique, on a souvent à disposition un n -échantillon tiré selon une distribution P . Le but est d'inférer des propriétés vérifiées par P à partir de propriétés satisfaites par sa version empirique P_n . Pour ce faire, d'une part il est important de montrer que les propriétés que l'on cherche à mettre en valeur sont préservées lorsque deux mesures sont proches au sens d'une certaine métrique. On parle alors de *stabilité*. D'autre part il est primordial que la métrique en question soit adaptée à l'échantillonnage. C'est-à-dire, qu'elle respecte la propriété de convergence faible de P_n vers P presque sûre (voir [Var58b] ou [BL14, Théorème 2.13]).

C'est le cas des métriques de Wasserstein, des distances adaptés à la convergence en loi et à l'échantillonnage que nous utilisons fréquemment dans cette thèse. En particulier, nous nous reposons sur les bornes en déviation et en espérance pour $W_1(P_n, P)$ obtenues dans [FG15] et [BL14]. Elles sont de l'ordre de $n^{-\frac{1}{d}}$ lorsque $d \geq 2$ et $n^{-\frac{1}{2}}$ lorsque $d = 1$.

¹¹On appelle *distribution empirique* associée à P , notée P_n , toute distribution uniforme sur un n -échantillon $\{X_1, X_2, \dots, X_n\}$ de loi P . C'est-à-dire, $P_n = \frac{1}{n} \sum_{i=1}^n \delta_{X_i}$.

¹²Une espace métrique (\mathcal{X}, δ) est un ensemble \mathcal{X} muni d'une application (*métrique* ou *distance*) qui est positive, vérifie l'hypothèse de séparation (à savoir, $\delta(x, y) = 0 \Rightarrow x = y$), symétrique et satisfaisant l'inégalité triangulaire.

¹³La tribu Borélienne associée à un espace métrique est la tribu engendrée par les ouverts.

Des métriques ou pseudo-métriques pour comparer des espaces métriques mesurés

Les métriques et autres outils introduits précédemment ont la particularité de mesurer une dissimilarité entre deux distributions définies sur le même espace. En particulier, on peut les utiliser pour comparer deux nuages de points générés selon deux distributions définies sur le même espace. Or, il arrive que l'on soit amené à comparer deux jeux de données qui ne sont pas inclus dans le même espace ou qui n'ont pas été mesurés dans le même système de coordonnées. Dans ce cas, il est plus facile de représenter les données comme un ensemble de points et d'associer une valeur à chaque paire de points ; la distance entre les deux points.

Le cadre mathématique que nous adoptons dans cette thèse et qui est parfaitement adapté à cette vision est celui des espaces métriques mesurés, étudiés par Gromov dans [GLP99]. Un *espace métrique mesuré* (ou mm-space) est défini comme un espace \mathcal{X} muni d'une métrique δ et d'une mesure de probabilité P définie sur la tribu des Boréliens.

De plus, on considère que deux espaces métriques mesurés (\mathcal{X}, δ, P) et (\mathcal{Y}, γ, Q) sont les mêmes lorsqu'ils sont *isomorphes*. C'est-à-dire, lorsqu'il existe une application $\phi: \mathcal{X} \rightarrow \mathcal{Y}$ satisfaisant les propriétés suivantes :

- ϕ est bijective,
- ϕ est une isométrie : $\gamma(\phi(x), \phi(x')) = \delta(x, x')$ pour tous $x, x' \in \mathcal{X}$,
- ϕ préserve la mesure : $Q(\phi(A)) = P(A)$ pour tout $A \in \sigma(\mathcal{X})$ (où $\sigma(\mathcal{X})$ est la tribu Borélienne engendrée par les ouverts de (\mathcal{X}, δ)).

Une telle application ϕ est appelée un *isomorphisme* entre les deux espaces métriques mesurés.

Dans ses travaux sur les espaces métriques mesurés [GLP99], Gromov offre une première caractérisation de tels espaces. En effet, dans son Théorème 3 $\frac{1}{2}$.5, il prouve que tout espace métrique mesuré peut être reconstruit à isomorphisme près, à partir de la connaissance pour toutes tailles d'échantillon n , de la loi de matrices de distances de taille $n \times n$ construites à partir de n -échantillons.

Plus récemment, dans [Mém11] Mémoli introduit des métriques définies sur l'ensemble des espaces métriques mesurés quotienté par la relation d'isomorphisme.¹⁴ Il s'agit des distances de Gromov-Wasserstein GW_p qui sont définies entre deux espaces métriques mesurés (\mathcal{X}, δ, P) et (\mathcal{Y}, γ, Q) pour un paramètre $p \geq 1$ par :

$$GW_p^p(P, Q) = \inf_{\pi \in \Pi(P, Q)} \frac{1}{2} \mathbb{E}_{(X', Y'), (X, Y) \sim \pi} [|\delta(X, X') - \gamma(Y, Y')|^p],$$

Ici, $\Pi(P, Q)$ est l'ensemble des plans de transport entre P et Q défini précédemment et (X, Y) et (X', Y') sont indépendants. Dans cette thèse, la distance de Gromov-Wasserstein de paramètre 1 GW_1 est notre métrique de référence lorsque nous cherchons à comparer des nuages de points ou des espaces métriques mesurés.

Malheureusement, même lorsque les espaces métriques mesurés sont discrets¹⁵, le calcul de la distance de Gromov-Wasserstein est extrêmement coûteux. Une vision alternative consiste à associer une *signature* à chaque espace métrique mesuré. Une signature est un objet mathématique par exemple un élément de \mathbb{R} , de \mathbb{R}^d ou une mesure de probabilité sur \mathbb{R} assigné à un espace métrique mesuré tel que la propriété suivante soit satisfaite : on assigne la même signature à deux espaces métriques mesurés lorsqu'ils sont isomorphes.

¹⁴Une métrique sur l'ensemble des espaces métriques mesurés \mathcal{MM} quotienté par la relation d'isomorphisme est une fonction $d: \mathcal{MM} \times \mathcal{MM} \rightarrow \mathbb{R}_+$ qui est symétrique et vérifie l'inégalité triangulaire. De plus, $d(M_1, M_2) = 0$ si et seulement si M_1 et M_2 sont isomorphes.

¹⁵Par *espace métrique discret*, on entend un espace métrique (\mathcal{X}, δ, P) où \mathcal{X} est fini.

Les espaces métriques mesurés peuvent ensuite être comparés via ces signatures. On notera $S(P)$ la signature associée à $(\mathcal{X}, \sigma(\mathcal{X}), P)$ et \mathcal{S} l'ensemble des signatures $S(P)$ pour tous les espaces métriques mesurés $(\mathcal{X}, \sigma(\mathcal{X}), P)$. Alors, toute distance $d_{\mathcal{S}}$ sur l'espace \mathcal{S} fournit une pseudo-distance $(P, Q) \mapsto d_{\mathcal{S}}(S(P), S(Q))$ sur l'espace des espaces métriques mesurés quotienté par la relation d'isomorphisme.

Un bon ensemble de signatures \mathcal{S} devrait satisfaire les propriétés suivantes :

- stabilité,
- discrimination voire séparation,
- facilité d'implémentation dans le cas discret.

La *stabilité* peut se caractériser par une borne du type $d_{\mathcal{S}}(S(P), S(Q)) \leq c GW(P, Q)$ pour $c > 0$. Le caractère *discriminant* est à comprendre dans le sens où lorsque $d_{\mathcal{S}}(S(P), S(Q))$ est petit, les espaces $(\mathcal{X}, \sigma(\mathcal{X}), P)$ et $(\mathcal{Y}, \sigma(\mathcal{Y}), Q)$ se ressemblent. En particulier, une bonne famille de signatures satisfait l'hypothèse de *séparation* au sens où $d_{\mathcal{S}}(S(P), S(Q)) = 0$ n'arrive que lorsque $(\mathcal{X}, \sigma(\mathcal{X}), P)$ et $(\mathcal{Y}, \sigma(\mathcal{Y}), Q)$ sont isomorphes. Enfin, une bonne signature devrait être facile à utiliser en pratique. C'est-à-dire, possible à calculer et à stocker lorsque P et Q sont des mesures discrètes, et telle que la distance $d_{\mathcal{S}}(S(P), S(Q))$ soit facile à calculer dans ce contexte.

C'est définitivement ce point de vue que nous adoptons dans cette thèse, lorsque nous cherchons à comparer deux nuages de points ou deux espaces métriques mesurés. Les signatures que nous introduisons sont des distributions sur \mathbb{R} et nous comparons deux signatures avec la métrique de Wasserstein W_1 .

L'exemple le plus basique de signatures qui soit consiste à assigner à tout espace métrique mesuré un élément quelconque Υ . Alors, deux espaces métrique mesurés isomorphes auront la même signature Υ . Ceci-dit, comme tous auront la même signature, cette signature n'a pas grand intérêt, elle ne permet pas de discriminer entre deux espaces métriques mesurés qui ne soient pas isomorphes. Un exemple simple de signature plus discriminant consiste à associer à $(\mathcal{X}, \mathcal{F}, \delta)$ le réel $\mathbb{E}_{X, X' \sim P}[\delta(X, X')]$ où X et X' sont indépendantes.

Dans [Mém11], Mémoli donne d'autres exemples de signatures. La distribution de forme (shape distribution) par exemple est définie comme la loi de $\delta(X, Y)$ pour X et Y deux variables aléatoires indépendantes de loi P . L'excentricité est la loi de $P\delta(X, u)$ pour $X \sim P$ et la distribution locale des distances est la distribution de $P(\mathcal{B}(X, r))$ pour un $r \geq 0$ fixé et $X \sim P$. Mémoli fournit des preuves de la stabilité de certaines d'entre elles par rapport à la distance de Gromov-Wasserstein GW_1 . Le caractère discriminant de la distribution de forme a été mis en avant dans [BK04]. Pourtant, un exemple dans [Blo77] montre que la condition de séparation n'est pas satisfaite. Mémoli fournit également un tel exemple pour la distribution locale des distances [Mém11, Exemple 5.6].

Les signatures sont fréquemment utilisées pour des tâches de classification ou simplement de pré-classification; voir par exemple [OFCD02]. D'un point de vue plus topologique, les diagrammes de persistance définis dans la Section 1.1.3 ont été utilisés à ces fins dans [CCSG⁺09, CDSO14]. Dans cette thèse, nous utilisons les signatures pour construire des tests statistiques permettant de comparer directement deux nuages de points.

Des tests statistiques pour comparer des nuages de points

Lorsque l'on a accès à deux mesures ou à deux espaces métriques mesurés, les métriques et pseudo-métriques introduites dans les paragraphes précédents permettent de décider directement si ce sont les mêmes mesures, si les espaces métriques mesurés sont les mêmes à isomorphisme près ou si elles/ils se ressemblent.

Dans le contexte des statistiques, les mesures et espaces métriques mesurés ne sont pas accessibles. Seuls des échantillons générés selon ces distributions sont à disposition. Du fait de l'aléa, deux échantillons tirés selon la même loi ne seront pas forcément identiques. Pourtant, on peut espérer que les lois empiriques associées soient proches. On ne pourra jamais avoir la certitude que les deux échantillons sont issus de la même loi. Cependant, il est possible dans certains cas de fournir une réponse à partir de ces données. On espère alors se tromper avec petite probabilité. Un procédé pour fournir une réponse aléatoire pertinente est donnée par la notion de test statistique.

Un *test statistique* est une variable aléatoire ϕ_n à valeurs dans $\{0, 1\}$. Plus précisément, ϕ_n est fonction d'une variable aléatoire T_n appelée *statistique de test*. Cette statistique de test est construite sur un échantillon de n points générés selon une procédure \mathcal{P}_θ dépendant d'un paramètre θ à valeurs dans un certain ensemble Θ . Le test ϕ_n est associé à deux hypothèses H_0 " $\theta \in \Theta_0$ " et H_1 " $\theta \in \Theta_1$ " où Θ_0 et Θ_1 sont deux sous ensembles de Θ disjoints. Idéalement, on aimerait que la variable aléatoire ϕ_n vaille 1 lorsque θ est dans Θ_1 et 0 lorsqu'il est dans Θ_0 .

La qualité d'un test statistique est mesurée en terme d'erreur de type I, qui est définie pour tout θ_0 dans Θ_0 par $\mathbb{P}_{\theta_0}(\phi_n = 1)$. Il s'agit de la probabilité de prétendre que θ est dans Θ_1 lorsque θ vaut en fait θ_0 . De plus, on dit qu'un test est de *niveau* $\alpha \in (0, 1)$ si son *erreur de type I* est majorée par α , c'est-à-dire si $\mathbb{P}_{\theta_0}(\phi_n = 1) \leq \alpha$ pour tous θ_0 dans Θ_0 . Deux tests statistiques de niveau fixé $\alpha \in (0, 1)$ peuvent être comparés à travers leur *erreur de type II*. Il s'agit de la fonction β définie pour tous θ_1 dans Θ_1 par $\mathbb{P}_{\theta_1}(\phi_n = 0)$, la probabilité de prétendre que θ est dans Θ_0 alors que θ vaut θ_1 . La *puissance* d'un test est la fonction définie par $1 - \beta$. Aussi, on dit qu'un test statistique ϕ_n est de *niveau asymptotique* α si pour tout θ_0 dans Θ_0 , $\mathbb{P}_{\theta_0}(\phi_n = 1)$ converge vers α lorsque la taille de l'échantillon n tend vers l'infini.

Les méthodes de construction de tests statistiques sont de plusieurs type. Ou bien la distribution de la statistique de test T_n ou sa loi limite (lorsque n tend vers l'infini) est connue sous l'hypothèse H_0 . Dans ce cas, l'expression du test ϕ_n en fonction de la statistique de test T_n se déduit directement de cette distribution sous H_0 . Ou bien la loi de T_n et la loi limite sont inconnues ou dépendent de P . Dans ce cas, une stratégie de type *bootstrap* ou de *sous-échantillonnage* permet parfois d'approcher la distribution de la statistique T_n sous l'hypothèse H_0 à partir des données. C'est le cas des tests statistiques que nous introduisons et étudions dans cette thèse.

Les méthodes de type *bootstrap* ont été introduites par Efron en 1979 ([Efr79]) dans le but de construire des intervalles de confiance, et ont été fréquemment utilisées depuis. Une référence pour le *bootstrap* avec un point de vue asymptotique est le livre de van der Vaart et Wellner [vdVW96]. Pour le point de vue non asymptotique, on pourra se référer à la thèse de Arlot [Arl07]. Afin de valider la stratégie d'approximation de la loi de T_n sous H_0 par une méthode de sous-échantillonnage, nous utilisons les métriques de Wasserstein. Dans le cadre des tests statistiques, une telle méthode de validation a déjà été employée dans [DBLL15] et [FLLRB12] par exemple.

Tester l'égalité de deux distributions définies sur un espace (Ω, \mathcal{F}) consiste à choisir pour Θ l'ensemble des couples de distributions (P, Q) et pour Θ_0 le sous ensemble des couples (P, P) .

Parmi les tests statistiques d'égalité présents dans la littérature, on peut citer le test de Kolmogorov-Smirnov pour des données générés dans \mathbb{R} . On peut également citer les tests de rangs-signés de Wilcoxon, les t-tests ou encore les tests de localisation et d'échelle. Des travaux récents sur le sujet comprennent les tests de Gretton et al. [GMBJR⁺12]. Ils reposent sur la dis-crépance moyenne maximale (MMD) dans les espaces de Hilbert à noyaux reproduisant (RKHS).

Dans cette thèse, nous construisons et étudions des tests d'égalité de deux espaces métriques mesurés à isomorphisme près. Ceci revient à considérer pour Θ l'ensemble des couples d'espaces métriques mesurés $((\mathcal{X}, \delta, P), (\mathcal{Y}, \gamma, Q))$ et pour Θ_0 un sous ensemble de Θ constitué des couples d'espaces métriques mesurés isomorphes :

$$\Theta_0 = \{((\mathcal{X}, \delta, P), (\mathcal{Y}, \gamma, Q)) \in \Theta \mid (\mathcal{X}, \delta, P) \text{ et } (\mathcal{Y}, \gamma, Q) \text{ sont isomorphes}\},$$

et

$$\Theta_1 = \Theta \setminus \Theta_0.$$

Jusqu'à présent, à notre connaissance, aucun tel test statistique n'a été étudié dans la littérature. On pourrait cependant imaginer un test de Kolmogorov-Smirnov entre un échantillon de la même loi que $\delta(X, X')$ où $X, X' \sim P$ sont indépendantes, et un autre échantillon de la loi de $\gamma(Y, Y')$ avec $Y, Y' \sim Q$ indépendantes. Dans cette thèse, nous comparons nos tests d'isomorphisme à cette procédure de test. Ils sont plus discriminants sous certaines alternatives. Une autre idée issue du domaine de la TDA consisterait à utiliser les intervalles de confiance calculés pour des outils topologiques tels que les diagrammes de persistance définis en section 1.1.3, dans [CFL⁺13] et [FLR⁺14]. Une telle méthode ne fonctionne pas car elle est trop conservatrice, le calcul des signatures est trop coûteux et un tel test ne tiendrait pas suffisamment compte de la mesure, seulement de son support.

Le domaine des tests statistiques d'égalité d'espaces métriques mesurés à isomorphisme près n'en est qu'à ses balbutiements.

1.1.2 Une introduction aux méthodes de partitionnement et de quantification

Diverses questions, par exemple celles des types en anthropologie, ou bien d'autres, d'ordre pratique, comme celles de la normalisation des objets industriels, exigent pour leur solution la détermination de n représentants fictifs d'une nombreuse population, choisis de manière à réduire autant que possible les écarts entre les éléments de la population et ceux de l'échantillon, l'écart étant mesuré entre tout élément réel et l'élément fictif qui lui est le plus proche.

Sur la division des corps matériels en parties, 1956 par H. STEINHAUS, [Ste56]

Si l'intérêt de faire des groupes à partir des données était déjà d'actualité en 1956, elle l'est toujours aujourd'hui, avec des applications dans les domaines des réseaux sociaux, des sciences sociales, de la psychologie, en imagerie médicale, en reconnaissance de forme, en exploration de données ou en apprentissage automatique... Les données actuelles étant massives, il est également capital de pouvoir réduire leur dimension sans pour autant perdre d'information.

Les réponses à ces questions portent le nom de *partitionnement*, lorsqu'il s'agit de faire des groupes à partir des données et de *quantification* lorsque l'objectif est de stocker l'information contenue dans un ensemble de données de façon compacte. Dans cette thèse, nous adoptons ces deux points de vue avec des méthodes assez similaires.

Le concept de quantification est né dans les années 1950' dans les laboratoires Bell. Une référence majeure en quantification est [GL00]. En 1957, Stuart Lloyd proposa un algorithme permettant de trouver une solution approchée à ce problème. La publication de ses travaux a eu lieu des années plus tard, en 1982, voir [Llo82]. En 1965, Forgy proposa dans [For65] un algorithme similaire à celui proposé par Lloyd. Le terme k -moyennes (k -means en anglais) est apparu dans le papier de MacQueen en 1967, [Mac67]. Ce concept repose sur les idées de Steinhaus dans [Ste56].

La méthode des k -moyennes correspond en fait exactement à la méthode de quantification telle qu'elle a été pensée par Lloyd, mais avec le point de vue k petit. En règle générale, toutes

les méthodes de partitionnement ou de quantification étudiées dans cette thèse reposent sur l'algorithme de Lloyd, de près ou de loin.

Méthode des k -moyennes et quantification – deux définitions équivalentes

On se place dans \mathbb{R}^d muni de la métrique Euclidienne $\|\cdot\|$. En quantification, on appelle *codebook* une famille $\mathbf{c} = (c_1, c_2, \dots, c_k)$ de k éléments (appelés *centres*) dans \mathbb{R}^d . On note $\mathbb{R}^{d,(k)}$ l'ensemble de tels codebooks. Étant donné un codebook \mathbf{c} , un *quantificateur* est une application $q : \mathbb{R}^d \rightarrow \{c_1, c_2, \dots, c_k\}$. Un quantificateur q^* est dit *optimal* au sens des k -moyennes lorsqu'il minimise le critère $q \rightarrow P\|u - q(u)\|^2$. De façon équivalente, un codebook \mathbf{c}^* est dit *optimal* au sens des k -moyennes lorsqu'il minimise le *critère de coût*

$$R : \mathbf{c} \mapsto P \min_{i \in \llbracket 1, k \rrbracket} \|u - c_i\|^2. \quad (1.3)$$

Le but de la quantification est de fournir une représentation compacte d'une mesure de probabilité P sur \mathbb{R}^d . Il s'agit de l'approcher au mieux, au sens de la métrique de Wasserstein W_2 , par une distribution dont le support fini contient au plus k points. On note $\mathcal{P}^{(k)}$ l'ensemble de ces distributions.

Le point de vue quantification et minimisation du critère des k -moyennes coïncident. En effet, en notant $P_{\mathbf{c}, \alpha} = \sum_{i=1}^k \alpha_i \delta_{c_i}$, avec $\mathbf{c} \in \mathbb{R}^{d,(k)}$ et des poids positifs α_i de somme 1, les éléments de $\mathcal{P}^{(k)}$, il vient :

$$W_2^2(P, \mathcal{P}^{(k)}) := \inf_{\mathbf{c}, \alpha} W_2^2(P, P_{\mathbf{c}, \alpha}) = \inf_{\mathbf{c}} P \min_{i \in \llbracket 1, k \rrbracket} \|u - c_i\|^2. \quad 16$$

La méthode de quantification ou de partitionnement des k -moyennes revient donc à déterminer la projection de P sur l'espace des distributions dont le support contient au plus k points, pour la métrique W_2 . Certaines distributions s'approchent facilement par des éléments de $\mathcal{P}^{(k)}$. C'est le cas des distributions à support compact qui sont (a, A, b) -Ahlfors-régulières¹⁷. Pour de telles distributions P , $W_2(P, \mathcal{P}^{(k)})$ est de l'ordre de $k^{-\frac{1}{b}}$; voir [GL00] ou [Klo12]. Lorsque \mathcal{N} est une distribution de moment d'ordre 2 fini et P est du type $P_{\mathbf{c}, \alpha} * \mathcal{N}$ ¹⁸, alors $W_2(P, \mathcal{P}^{(k)})$ est majorée par $\mathcal{N}\|u\|^2$. Cette majoration est en fait une égalité lorsque les supports des mesures $\delta_{c_i} * \mathcal{N}$ sont disjoints.

Décomposition de l'espace en cellules de Voronoï et Algorithme de Lloyd

La méthode des k -moyennes repose sur la décomposition de l'espace en cellules de Voronoï. La *cellule de Voronoï* V_i associée à un centre c_i dans un codebook \mathbf{c} est définie par :

$$V_i = \left\{ x \in \mathbb{R}^d \mid \|x - c_i\| \leq \|x - c_j\|, \forall j \neq i \right\}.$$

On peut également définir les cellules $V_i^\circ = \{x \in \mathbb{R}^d \mid \|x - c_i\| < \|x - c_j\|, \forall j \neq i\}$. Dans la suite, nous travaillerons avec des cellules W_i telles que :

¹⁶On remarquera que le carré de la distance de Wasserstein $W_2^2(P, P_{\mathbf{c}, \alpha})$ se réécrit comme $P\|u - q_{\mathbf{c}, \alpha}(u)\|^2$ où $q_{\mathbf{c}, \alpha}$ est un quantificateur à valeurs dans \mathbf{c} naturellement associé au plan de transport optimal entre P et $P_{\mathbf{c}, \alpha}$. En particulier, $W_2^2(P, P_{\mathbf{c}, \alpha}) \geq P \min_{i \in \llbracket 1, k \rrbracket} \|u - c_i\|^2$. Rappelons que la distance $W_2(P, Q)$ s'écrit comme $\inf_{\pi \in \Pi(P, Q)} \mathbb{E}_{(X, Y) \sim \pi} [\|X - Y\|^2]$ et qu'un *plan de transport optimal* est un plan de transport π^* tel que $W_2(P, Q) = \mathbb{E}_{(X, Y) \sim \pi^*} [\|X - Y\|^2]$.

¹⁷Une distribution P est dite (a, A, b) -Ahlfors-régulière lorsqu'il existe a, A et b positifs tels que : $\forall x \in \text{Supp}(P)$, $1 \wedge ar^b \leq P(\mathcal{B}(x, r)) \leq Ar^b$

¹⁸La mesure de probabilité $P * Q$ correspond à la distribution d'une variable aléatoire $X + Y$ lorsque $X \sim P$, $Y \sim Q$ et X et Y sont indépendantes.

- $V_i^\circ \subset W_i \subset V_i$ pour tout i ,
- La famille $(W_i)_{i \in \llbracket 1, k \rrbracket}$ forme une *partition*¹⁹ de \mathbb{R}^d .

Afin de comprendre le lien entre la décomposition d'un espace en cellules de Voronoï et la méthode des k -moyennes, remarquons que

$$R(\mathbf{c}) = P \min_{i \in \llbracket 1, k \rrbracket} \|u - c_i\|^2 = P \sum_{i=1}^k \mathbb{1}_{W_i}(u) \|u - c_i\|^2.$$

Étant donné un codebook optimal $\mathbf{c}^* = (c_1^*, c_2^*, \dots, c_k^*)$ et W_i^* les cellules de Voronoï associées, tout quantificateur qui envoie les éléments de W_i^* sur c_i^* est optimal. De plus, étant donnée une partition de l'espace $(W_i)_{i \in \llbracket 1, k \rrbracket}$, les centres minimisant $P \sum_{i=1}^k \mathbb{1}_{W_i}(u) \|u - c_i\|^2$ sont donnés par les espérances $c_i = |W_i|^{-1} P u \mathbb{1}_{W_i}(u)$.

L'algorithme de Lloyd illustré par la Figure 1.1 repose sur ces deux propriétés. Afin d'initialiser l'algorithme, on se donne des centres c_1, c_2, \dots, c_k au hasard parmi les données. Pour une initialisation plus adaptée, on pourra se référer à l'algorithme k -means++ [AV07]. Ensuite, on répète les deux étapes suivantes jusqu'à ce que l'algorithme converge, c'est-à-dire jusqu'à ce que les centres c_1, c_2, \dots, c_k ne varient plus.

- *Étape 1* : Décomposition de \mathbb{R}^d en cellules de Voronoï associées aux centres $c_1, c_2, \dots, c_k : W_1, W_2, \dots, W_k$.
- *Étape 2* : Calcul des nouveaux centres donnés pour tout i par $c_i = |W_i|^{-1} P u \mathbb{1}_{W_i}(u)$.

Concrètement, lorsque P est une mesure uniforme sur un ensemble fini de n points \mathbb{X}_n , la première étape consiste à faire une partition de \mathbb{X}_n en k paquets C_1, C_2, \dots, C_k où les éléments de C_i sont dans W_i . La seconde étape consiste à actualiser les centres c_i . Le nouveau centre c_i est donné par la moyennes des éléments du i -ème paquet C_i .

Cette méthode permet de faire k groupes à partir de n'importe quel nuage de points dans \mathbb{R}^d . Aucune information supplémentaire n'est requise. En ce sens, la méthode des k -moyennes est une méthode de *partitionnement non supervisé*.

La méthode des k -moyennes avec l'algorithme de Lloyd est adaptée à des mesures P qui se décomposent en k mesures bien séparées, de préférence isotropes et de même variance. C'est le cas des mélanges Gaussiens de centres c_1, c_2, \dots, c_k , de même variance σ^2 . Dans ce cas, la cellule de Voronoï associée à c_i coïncide avec la région de l'espace où la densité de $\mathcal{N}(c_i, \sigma^2)$ est supérieure aux densités associées aux autres centres c_j .

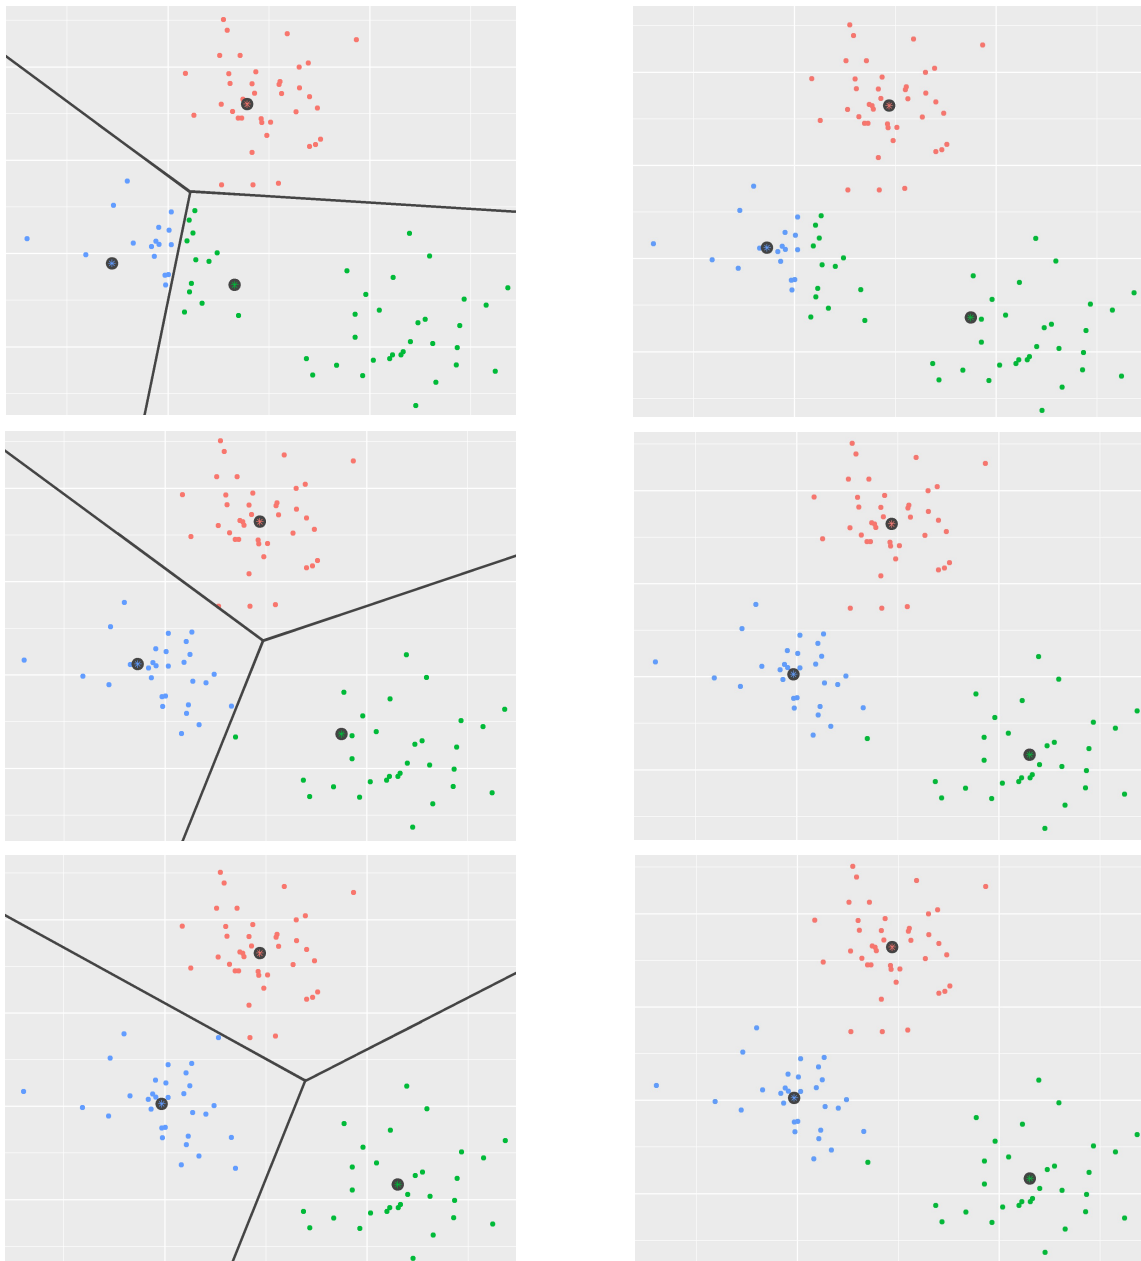
Ces méthodes s'adaptent à des mélanges non isotropes.

Adaptation de la méthode des k -moyennes pour les divergences de Bregman

Lorsque le modèle de lois considérées n'est pas isotrope et est hétéroscédastique, au sens où les variances des lois du mélange diffèrent, la vraisemblance crée une distorsion de l'espace et la décomposition en cellules de Voronoï n'est plus adaptée au problème de partitionnement. C'est le cas de la famille des lois de Poisson $(\mathcal{P}(c))_{c \geq 0}$ par exemple. La vraisemblance associée à une loi $\mathcal{P}(c)$ étant donné $x \in \mathbb{N}$ s'écrit

$$f_c(x) = \frac{c^x}{x!} \exp(-c) = \exp(-d_\phi(x, c) + C(x)),$$

¹⁹La famille $(W_i)_{i \in \llbracket 1, k \rrbracket}$ forme une partition de \mathbb{R}^d si $\bigcup_{i \in \llbracket 1, k \rrbracket} W_i = \mathbb{R}^d$ et si les W_i sont deux à deux disjoints.

FIGURE 1.1 : Algorithme de Lloyd pour la méthode des k -moyennes

avec $C(x)$ une fonction de x et $d_\phi(x, c) = x \log\left(\frac{x}{c}\right) - (x - c)$ la divergence de Bregman associée à la fonction convexe $\phi : x \mapsto x \log(x) - x$.

La *divergence de Bregman* associée à une fonction convexe ϕ différentiable sur $\Omega \subset \mathbb{R}^d$, de gradient $\nabla\phi$ est définie dans [Bre67] pour tout x et y de Ω par

$$d_\phi(x, y) = \phi(x) - \phi(y) - \langle \nabla_y \phi, x - y \rangle.$$

Une divergence de Bregman n'est pas une métrique au sens où l'inégalité triangulaire et la symétrie ne sont pas nécessairement vérifiées. Cependant, d_ϕ est toujours positive et lorsque de plus ϕ est strictement convexe, d_ϕ vérifie l'hypothèse de séparation au sens où $d_\phi(x, y) = 0$ si et seulement si $x = y$.

En choisissant pour fonction ϕ le carré de la norme Euclidienne $\|\cdot\|^2$, la divergence de Bregman coïncide avec le carré de la distance Euclidienne, $d_\phi(x, y) = \|x - y\|^2$.

On peut définir les cellules de Bregman-Voronoi associées à une divergence de Bregman d_ϕ et à un codebook c en remplaçant le carré de la norme Euclidienne par d_ϕ dans la définition des cellules de Voronoi :

$$V_i = \left\{ x \in \mathbb{R}^d \mid d_\phi(x, c_i) \leq d_\phi(x, c_j), \forall j \neq i \right\}.$$

Une étude poussée des réseaux de Bregman-Voronoi est disponible dans [BNN10] et [NBN07].

Dans le cadre du modèle des lois de Poisson, la partition de l'espace en cellules de Bregman-Voronoi pour la divergence associée à la fonction convexe $\phi : x \mapsto x \log(x) - x$ est tout à fait adaptée. En effet, un point x est dans la cellule de Bregman-Voronoi V_i lorsque la vraisemblance $c \rightarrow f_c(x)$ est minimale en c_i parmi les éléments de $\{c_1, c_2, \dots, c_k\}$.

Plus généralement, le lien entre les modèles de mélange et les divergences de Bregman a été établi dans [BMDG05], dans le cadre du partitionnement de type k -moyennes. L'étude comprend les mélanges de lois gaussiennes, de Poisson, binomiales etc.

L'algorithme de Lloyd s'adapte parfaitement aux divergences de Bregman. Les étapes de l'algorithme sont les suivantes :

- *Étape 1* : Décomposition de \mathbb{R}^d en cellules de Bregman-Voronoi associées aux centres $c_1, c_2, \dots, c_k : W_1, W_2, \dots, W_k$.
- *Étape 2* : Calcul des nouveaux centres donnés pour tout i par $c_i = \frac{1}{|W_i|} \int_{W_i} u \, dQ(u)$.

On peut remarquer que la seconde étape de l'algorithme n'a pas été modifiée. En effet, pour toute mesure de probabilité Q , l'application $c \mapsto Qd_\phi(u, c)$ atteint son minimum en l'espérance, $c = Qu$.²⁰

Dans [BGW05], Banerjee montre que les seules fonctionnelles F pour lesquelles $QF(u, c)$ est minimale en $c = Qu$ pour toute distribution Q sont de la forme $F = F_0 + C$ où F_0 est une divergence de Bregman et C une constante. L'algorithme de Lloyd est alors adapté aux divergences de Bregman mais n'est a priori pas adapté à d'autres types de fonctionnelles. Par exemple, dans [CAGM97], Cuesta-Albertos et al. remplacent le carré de la norme Euclidienne par une fonction convexe de celle-ci dans la méthode des k -moyennes. Une telle fonctionnelle n'est pas une divergence de Bregman. Ils n'ont pas pu utiliser l'algorithme de Lloyd mais ont eu recours à un algorithme de type Metropolis Hasting.

L'adaptation de la méthode des k -moyennes aux divergences de Bregman par Banerjee et al. dans [BMDG05] fournit un ensemble de méthodes de partitionnement ou de quantification. De façon générale, il est possible d'aggréger différentes méthodes de partitionnement et d'en mesurer la qualité en termes d'information mutuelle ; voir [SJ02].

Dans cette thèse, nous adaptons cette méthode de partitionnement Bregman à des données corrompues par des données aberrantes, en trimmant.

²⁰ On le voit facilement en notant d'une part que $Qd_\phi(u, c) = Qd_\phi(u, Qu) + d_\phi(Qu, c)$, et d'autre part que $d_\phi(Qu, c)$ est positive et nulle seulement lorsque $c = Qu$.

Méthodes des k -moyennes trimmées

Dans [CAGM97], Cuesta-Albertos et al. adaptent la méthode des k -moyennes à des distributions de type $Q = \alpha P + (1 - \alpha)R$ ²¹ avec $\alpha \in (0, 1]$, P une distribution de type mélange Gaussien et R une distribution quelconque. Un échantillon tiré selon Q est composé de k groupes de points avec des données aberrantes issues de la distribution R .

Étant donné un paramètre de masse $h \in (0, 1]$, la méthode proposée consiste à séparer la mesure Q en deux mesures Q_{signal} et Q_{bruit} de masses respectives h et $1 - h$. Il s'agit d'appliquer à la distribution Q_{signal} la méthode des k -moyennes, c'est-à-dire, de minimiser le critère

$$(\mathbf{c}, Q_h) \rightarrow Q_h \min_{i \in [1, k]} \|u - c_i\|^2$$

en \mathbf{c} , un codebook et Q_h , une mesure de probabilité telle que $hQ_h = Q_{signal}$ soit une sous-mesure de Q . De façon équivalente, cela revient à minimiser le risque $R(\mathbf{c})$ en \mathbf{c} , où :

$$R(\mathbf{c}) = \inf_{Q_h} Q_h \min_{i \in [1, k]} \|u - c_i\|^2.$$

Dans ce papier, les auteurs démontrent l'existence d'un codebook optimal \mathbf{c}^* . La mesure Q_h optimale correspond alors à la restriction de Q à une union de k boules de même rayon, centrées en les éléments de \mathbf{c}^* . Soit $\hat{\mathbf{c}}_n$ le codebook optimal associé à la distribution Q_n . Lorsque \mathbf{c}^* est unique, Cuesta-Albertos et al. généralisent les résultats de convergence presque sûre de $\hat{\mathbf{c}}_n$ vers \mathbf{c}^* de Pollard, [Pol82b].

Ils démontrent également la convergence presque sûre du risque $R(\hat{\mathbf{c}}_n)$ vers $R(\mathbf{c}^*)$. Dans cette thèse, nous précisons ces résultats avec une étude non asymptotique de cette convergence.

La méthode de k -moyennes trimés est adaptée aux mélanges de lois de même variance. Dans le cas des modèles Gaussiens hétéroscédastiques, des méthodes de partitionnement trimmé ont également été proposées, voir par exemple la méthode de [GEGMMI08] implémentée par l'algorithme R tclust. Cette méthode repose sur la vraisemblance trimmée introduite par Neykov et Neytchev en 1990 [NN90]. En général, les méthodes permettant de traiter ce type de données reposent sur l'algorithme EM²² (Expectation Maximisation).

Dans cette thèse, nous proposons et étudions une autre méthode implémentée par un algorithme de type Lloyd trimmé.

Résultats d'approximation des k -moyennes à partir de données

Soit P_n la mesure empirique associée à un n -échantillon de loi P . D'après Pollard [Pol82b], la méthode de k -moyennes est consistante au sens où sous la condition d'existence et d'unicité d'un codebook optimal \mathbf{c}^* pour P , toute suite $\hat{\mathbf{c}}_n$ de codebooks optimaux pour P_n converge vers ce codebook \mathbf{c}^* presque sûrement. De plus, sous certaines hypothèses supplémentaires, Pollard montre dans [Pol82a] que les codebooks optimaux $\hat{\mathbf{c}}_n$ vérifient un théorème central limite.

Les premiers résultats non asymptotiques pour la méthode des k -moyennes dans \mathbb{R}^d apparaissent dans le papier de Linder [Lin02]. Les bornes obtenues sont du type :

$$\mathbb{E} [R(\hat{\mathbf{c}}_n) - R(\mathbf{c}^*)] \leq CK^2 \sqrt{\frac{kd}{n}},$$

²¹Une variable aléatoire X suit la loi $\alpha P + (1 - \alpha)R$ lorsqu'avec probabilité α , $X \sim P$ et avec probabilité $1 - \alpha$, $X \sim R$.

²²Cet algorithme introduit par Redner et Walker en 1984 ([RW84], voir [FR02] pour une référence plus récente) repose sur la notion de vraisemblance.

lorsque le support de P est inclus dans la boule Euclidienne $\mathcal{B}(0, K)$ avec C une constante universelle. De telles vitesses lentes peuvent être obtenues avec la théorie de Vapnik [Vap82], voir aussi la version plus récente de Lugosi [Lug02].

Dans le cadre plus général des espaces de Hilbert, Biau et al. prouvent dans [BDL08] des bornes en $\frac{k}{\sqrt{n}}$, qui ne dépendent plus de la dimension. Avec des techniques similaires, Fischer démontre dans [Fis10] des bornes en $\frac{k}{\sqrt{n}}$ pour l'excès de risque lorsque $\|\cdot\|^2$ est remplacée par une divergence de Bregman.

Des vitesses plus rapides pour la méthode des k -moyennes dans \mathbb{R}^d en $\frac{1}{n}$ et sous des hypothèses de condition de marge ont été obtenues dans [Lev15], voir aussi la thèse de Levrard [Lev14]. Ces vitesses reposent par exemple sur les travaux de Massart [Mas07, Théorème 5.1].

1.1.3 Une courte introduction à l'inférence géométrique et topologique

Inférence topologique avec la fonction distance au compact

Dans les domaines de l'inférence topologique et géométrique, le but est d'approcher la topologie ou la géométrie d'un objet \mathcal{K} (par exemple une sous-variété de \mathbb{R}^d) à partir d'un nuage de points échantillonnés dans un voisinage de cet objet. Les méthodes fréquemment utilisées reposent sur l'étude de sous-niveaux de fonctions distance.

Étant donné un espace métrique (\mathcal{X}, δ) et un compact $\mathcal{K} \subset \mathcal{X}$, la *distance au compact* \mathcal{K} est la fonction $d_{\mathcal{K}} : \mathcal{X} \rightarrow \mathbb{R}_+$, définie pour tout $x \in \mathcal{X}$ par $d_{\mathcal{K}}(x) = \inf_{y \in \mathcal{K}} \delta(x, y)$. Les sous-niveaux $\mathcal{K}^\epsilon = d_{\mathcal{K}}^{-1}(\epsilon)$ de la fonction $d_{\mathcal{K}}$ correspondent aux ϵ -épaississements²³ de \mathcal{K} . La *distance de Hausdorff* d_H définie par $d_H(\mathcal{K}, \mathcal{K}') := \|d_{\mathcal{K}} - d_{\mathcal{K}'}\|_\infty$ est fréquemment utilisée en analyse topologique ou géométrique des données pour comparer la topologie de deux compacts, voir par exemple la thèse de Aamari [Aam17].

La topologie des données peut par exemple être décrite par des diagrammes de persistance²⁴ (introduits en 2002 par Edelsbrunner et al. dans [ELZ02]) ou des code-barres²⁵. Ces deux outils décrivent l'évolution de l'homologie (nombres de composantes connexes, nombres de trous etc.) des sous-niveaux \mathcal{K}^ϵ d'une fonction distance, lorsque ϵ va de 0 à l'infini. Dans la Figure 1.2, nous avons tracé la fonction distance à un ensemble de 4 points, ainsi que le code-barre associé.

L'information contenue dans un diagramme de persistance peut également être résumée dans un code-barre.

Deux diagrammes de persistance $D(f)$ et $D(g)$ associés aux sous-niveaux respectifs de fonctions f et g suffisamment régulières sont proches au sens de la distance de Bottleneck²⁶ entre diagrammes dès lors que f et g sont proches. En effet, sous certaines hypothèses de régularité

²³Le ϵ -épaississement \mathcal{K}^ϵ est défini par $\mathcal{K}^\epsilon = \{x \in \mathcal{X} \mid \exists y \in \mathcal{K}, \delta(x, y) \leq \epsilon\}$.

²⁴Un *diagramme de persistance* est une famille de paires de points (b, d) . Le temps de naissance b_i correspond au paramètre ϵ pour lequel la i -ème composante (e.g. composante connexe, trou...) apparaît. Le temps de mort d_i correspond au paramètre ϵ pour lequel la i -ème composante disparaît. Par exemple, lorsque l'on considère le nombre de composantes connexes d'un ensemble de n points \mathcal{K}_n , le diagramme de persistance contient exactement n paires (b, d) . Toutes ces paires satisfont $b = 0$ puisque $\mathcal{K}_n^0 = \mathcal{K}_n$ est composé de n composantes connexes. Aussi, $(0, \infty)$ sera dans le diagramme puisque lorsque ϵ est suffisamment grand, \mathcal{K}_n^ϵ contient exactement une composante connexe.

²⁵Un code-barre est constitué de l'ensemble des segments $[b_i, d_i]$ pour les points (b_i, d_i) contenus dans le diagramme de persistance.

²⁶La distance de bottleneck est définie par $dB(D, D') = \inf_\rho \sup_{x \in D} \|x - \rho(x)\|_\infty$, avec ρ la fonction bijective entre les points x du diagramme D et les points du diagramme D' , avec éventuellement des points supplémentaires $x = (b, b)$ sur la diagonale dans D et $\rho(x)$ sur la diagonale dans D' .




FIGURE 1.2 : Distance to a set of four points and barcode

sur les fonctions f et g , on peut écrire :




$$dB(D(f), D(g)) \leq \|f - g\|_\infty. \quad (1.4)$$

Il s'agit du théorème de stabilité des diagrammes de persistance par [CSEH07]. Lorsque les fonctions en question sont des fonctions distance à des compacts \mathcal{K} et \mathcal{K}' , $d_{\mathcal{K}}$ et $d_{\mathcal{K}'}$, alors cette distance est majorée par la distance de Hausdorff $d_H(\mathcal{K}, \mathcal{K}')$. Par conséquent, la distance de Hausdorff est particulièrement adaptée à l'inférence topologique des données.

Dans cette thèse, nous adoptons un point de vue adapté à la présence de données aberrantes. Nous travaillons avec une généralisation de la distance au compact, la distance à la mesure.

Inférence topologique avec les fonctions distance à la mesure

Un problème majeur de la distance de Hausdorff est qu'elle n'est pas stable en présence de données aberrantes. En effet, l'ajout de quelques points à un ensemble compact \mathcal{K} peut induire une forte variation de la distance au compact \mathcal{K} . Alors, les diagrammes de persistance ou les code-barres associés changent de façon drastique. Cette difficulté est illustrée par les Figures 1.3, 1.4 puis 1.6.

FIGURE 1.3 : Distance à \mathcal{K} FIGURE 1.4 : Distance à \mathcal{K}_n FIGURE 1.5 : DTM avec $h = 0.2$

Afin de remédier à ce problème, la distance à la mesure a été introduite par Chazal, Cohen-Steiner et Mérigot en 2009. Elle est définie comme une généralisation de la distance au compact, pour une distribution P et un paramètre de masse $h \in (0, 1]$. Elle s'exprime à partir de la pseudo-distance $\delta_{P,h}$ définie pour tout $l \in [0, 1]$ par

$$\delta_{P,l}(x) = \inf\{r > 0 \mid P(\overline{\mathcal{B}}(x, r)) > l\},$$

où $\overline{\mathcal{B}}(x, r)$ désigne la boule Euclidienne fermée, centrée en x et de rayon r . La distance à la mesure (DTM) $d_{P,h}$ est alors définie par

$$d_{P,h}^2(x) = \frac{1}{h} \int_{l=0}^h \delta_{P,l}^2(x) dl$$




FIGURE 1.6 : Code-barres associés aux figures 1.3, 1.4 et 1.5

ou plus généralement, pour $r \geq 1$ par

$$d_{P,h,r}(x) = \left(\frac{1}{h} \int_{l=0}^h \delta_{P,l}^r(x) dl \right)^{\frac{1}{r}}$$

Lorsque la mesure P est une mesure uniforme sur un ensemble fini de points \mathbb{X}_n , la puissance r -ème de la distance à la mesure en un point x s'exprime comme la moyenne des puissances r -mes des distances entre x et ses $q = hn$ plus proches voisins dans \mathbb{X}_n . Elle est donc facilement implémentable en pratique.

Lorsque $r = 2$, une autre définition de la distance à la mesure en termes de sous-mesures est donnée dans [CCSM11], par :

$$d_{P,h,2}^2(x) = P_{x,h} \|u - x\|^2,$$

où $P_{x,h}$ est la restriction de P à la boule de centre x et de masse h pour P . En règle générale,

$$d_{P,h,2}^2(x) = \inf_{\tilde{P}} \tilde{P} \|u - x\|^2.$$

L'infimum est pris sur l'ensemble des distributions \tilde{P} telles que $h\tilde{P}$ soit une sous-mesure de P de masse h . Cette dernière définition permet d'exprimer la DTM comme une fonction puissance :

$$d_{P,h,2}^2(x) = \inf_{\tilde{P}} \|m(\tilde{P}) - x\|^2 + v(\tilde{P}),$$

où $m(\tilde{P})$ est l'espérance de \tilde{P} et $v(\tilde{P})$ sa variance.

Lorsque $r = 2$, la fonction $\psi_{P,h} : x \mapsto \|x\|^2 - d_{P,h}^2(x)$ est convexe. Dans cette thèse, la méthode de partitionnement/quantification Bregman de [BMDG05] appliqué à cette fonction convexe. Une version robuste (aux données aberrantes) de la méthode des k -moyennes. Nous étudions une telle méthode.

La distance à la mesure est stable par rapport à la distance de Wasserstein. Voir [CCSM11] pour le cas de \mathbb{R}^d et [Buc14] pour le cas plus général des espaces métriques. Lorsque P et Q sont deux mesures de probabilité supportées sur le même espace (\mathcal{X}, δ) , on a

$$\|d_{P,h,r} - d_{Q,h,r}\|_{\infty} \leq \frac{1}{h^{\frac{1}{r}}} W_r(P, Q).$$

La norme $\|\cdot\|_{\infty}$ correspond au suprémum pris sur l'ensemble des éléments x de \mathcal{X} .

Cette borne en Wasserstein permet d'obtenir des vitesses de convergence de l'ordre de $n^{-\frac{1}{2}}$ pour la déviation $\|d_{P,h} - d_{P_n,h}\|_{\infty}$. Cette vitesse peut être améliorée lorsque le paramètre de masse h vu comme une fonction de n décroît vers 0; voir [CMM16].

Le théorème de stabilité des diagrammes de persistance (1.4) s'applique aux fonctions de type DTM. En particulier, deux mesures proches au sens de Wasserstein auront sensiblement les

mêmes diagrammes de persistance (associés à la DTM). La DTM est une bonne approximation de la distance au support pourvu que h soit suffisamment petit [CCSM11]. Donc, l'information topologique contenue dans des données bruitées peut être lue dans le diagramme de persistance associé à la distance à la mesure.

D'un point de vue pratique, calculer un diagramme de persistance associé à la distance à un ensemble de n points revient à calculer l'homologie d'une union de n boules, pour différents rayons. Un tel calcul est accessible à condition que n ne soit pas trop élevé.

Les sous-niveaux de la DTM associée à un nuage de n points sont des unions d'environ $\binom{n}{q}$ boules, avec $q = hn$ le nombre de plus proches voisins considérés dans le calcul de la distance à la mesure. Le calcul d'invariants topologiques est alors compliqué. Afin de rendre ce calcul possible, des approximations de la distance à la mesure dont les sous-niveaux sont des unions de n boules ont été introduites dans [GMM11] (la k -witnessed distance) et dans [BCOS15] (autres fonctions puissance).

Dans la littérature, à notre connaissance, il n'existe pas encore d'approximation de la distance à la mesure dont les sous-niveaux soient des unions de k boules, avec k sous-linéaire par rapport à la taille des données n . Dans cette thèse, nous proposons et étudions une telle approximation de la distance à la mesure. Ceci rend possible le calcul d'invariants topologiques pour des données massives et bruitées.

Pour aller plus loin. L'intérêt des méthodes d'inférence topologique va au delà du domaine de l'analyse topologique des données. En effet, les méthodes de partitionnement hiérarchique telles que la méthode de *single linkage* (initialement introduite en 1951 par Florek et al. dans [FLP+51] puis dans [G.65] et [Joh67]), sont fortement liées à ces outils topologiques. Cette méthode est également connue sous le nom de partitionnement des k plus proches voisins (nearest neighbour clustering). Une telle méthode fournit une famille de partitionnements des données, indexée par un paramètre $\alpha \in [0, +\infty]$. Le partitionnement d'un nuage de n points \mathcal{K}_n associé au paramètre α correspond à l'ensemble des groupes constitués de points appartenant à la même composante connexe dans \mathcal{K}_n^α . L'évolution des composantes est donnée par le diagramme de persistance associé à la distance à \mathcal{K} .

On pourrait imaginer une procédure de type single-linkage associée aux sous-niveaux de la DTM. Une telle méthode permettrait de palier le principal défaut de la méthode de single-linkage, son caractère non robuste aux données aberrantes. C'est plus ou moins l'idée des travaux de Chazal et al. (voir [CGOS13]) où une méthode de partitionnement reposant sur les diagrammes de persistance est étudiée.

1.2 Contributions et Contenu

Dans cette thèse, nous traitons trois types de questions. La première concerne la *comparaison de deux nuages de points*. Notre contribution consiste à introduire et étudier un test statistique permettant de décider si deux jeux de données sont issus de deux espaces métriques mesurés égaux, à isomorphisme près. La seconde question concerne la *décomposition d'un nuage de points corrompu par des données aberrantes en plusieurs groupes pertinents*. D'une part, nous étudions la généralisation de la méthode de partitionnement des k -moyennes trimmées aux divergences de Bregman. Cette méthode permet de traiter le cas des modèles de mélange de lois Gaussiennes homoscédastiques, binomiales, Gamma, de Poisson etc. D'autre part, nous développons et étudions une nouvelle méthode pour partitionner des données échantillonnées selon un modèle Gaussien hétéroscédastique, en présence de données aberrantes. La dernière question concerne

l'approximation d'un ensemble compact par un ensemble de k centres, à partir d'un nuage de points tirés dans un voisinage de ce compact, avec d'éventuelles données aberrantes. Nous développons pour ce faire une nouvelle méthode de quantification basée sur la fonction distance à la mesure.

Les deux dernières questions sont fortement liées. Aussi, nous appliquons des méthodes similaires dans nos travaux afin de les traiter. De plus, dans cette thèse, nous adoptons un point de vue géométrique. En effet, les résultats et méthodes développées reposent fortement sur l'outil géométrique *distance à la mesure*.

Étant donnés deux nuages de points, une question importante en analyse de données est de savoir si les distributions dont ils sont issus se ressemblent. Par "se ressembler", on peut entendre "être égales". Dans la littérature, de nombreux tests statistiques ont été développés pour répondre à cette question particulière. On pense aux tests de Kolmogorov-Smirnov et plus récemment aux tests de Gretton et al. dans [GMBJR⁺12].

Dans cette thèse, nous cherchons à savoir si les distributions sous-jacentes sont égales modulo une certaine transformation rigide (e.g. modulo une translation, rotation, symétrie...). Plus précisément, nous construisons une famille de tests statistiques permettant de décider si les espaces métriques mesurés sous-jacents sont égaux à isomorphisme près. À notre connaissance, aucun tel test statistique n'a été étudié dans la littérature. Ce nouveau point de vue permet de comparer des données qui n'ont pas été mesurées dans le même espace ou dans le même système de coordonnées. Seule la connaissance des distances entre couples de points au sein d'un même échantillon est requise pour la construction de nos tests.


Pour construire de tels tests statistiques, nous procédons comme suit. Nous décidons de rejeter l'hypothèse d'égalité des espaces métriques mesurés à isomorphisme près lorsque les deux nuages de points sont trop différents. C'est-à-dire, lorsqu'une certaine "distance" entre les deux distributions empiriques associées est trop grande ; plus grande qu'une certaine valeur critique. Un bon test statistique consiste alors à choisir une bonne "distance" et une bonne *valeur critique*.

Dans cette thèse, nous comparons deux espaces métriques mesurés à isomorphisme près par le biais de la distance de Gromov-Wasserstein GW_1 . Cependant, comme mentionné précédemment, une telle distance n'est pas calculable en pratique, même dans le cas discret. C'est pourquoi nous adoptons le point de vue des signatures.

Nous introduisons une famille de signatures faciles à calculer en pratique. Leur calcul revient à une recherche de plus proches voisins. Nous montrons que les pseudo-métriques induites par ces familles de signatures sont stables par rapport à GW_1 et discriminantes sous certaines alternatives. Ces propriétés assurent que les tests soient du bon niveau et puissants.

La valeur critique est à déterminer en fonction du comportement de la pseudo-distance entre deux échantillons, sous l'hypothèse H_0 . Cependant, ce comportement dépend des distributions sous-jacentes. Afin de remédier à cette difficulté, nous développons une méthode de type sous-échantillonnage. Cette méthode permet de mimer le comportement de la pseudo-distance sous l'hypothèse H_0 . On en déduit une valeur critique, construite à partir des données.

Nous validons notre procédure d'échantillonnage à l'aide de la distance de Wasserstein W_1 . L'utilisation des distances de Wasserstein dans ce contexte est encore peu répandue dans la littérature ; voir [DBLL15, FLLRB12].



Un test de Kolmogorov-Smirnov entre deux échantillons de paires de distances serait une alternative envisageable à notre famille de tests statistiques. Cependant, nos tests sont plus discriminants dans certains cas. Par exemple, ils permettent de discriminer entre les deux spirales ci-contre, qui ne sont pas isomorphes. Le test de Kolmogorov-Smirnov ne le permet pas.

Étant donné un nuage de points échantillonnés ou bien selon un mélange de k distributions, ou bien dans le voisinage d'un ensemble compact (e.g. une sous-variété), en présence de données aberrantes, on se demande comment le réduire à un ensemble de k centres. Dans un premier temps, on prend le point de vue du partitionnement des données (i.e. k petit). Dans un second temps, on prend le point de vue de la quantification (i.e. k grand). Les méthodes proposées reposent sur des généralisations de la méthode des k -moyennes et de l'algorithme de Lloyd.

Dans cette thèse, nous proposons des méthodes qui permettent de rejeter d'éventuelles données aberrantes. Pour ce faire, nous adoptons le point de vue des k -moyennes trimmées de Cuesta-Albertos et al. ; voir [CAGM97]. Contrairement aux méthodes usuelles, cela ne consiste pas à éliminer les données aberrantes une fois les k centres optimaux déterminés, mais à en tenir compte à chaque étape de l'algorithme de recherche de ces centres.²⁷

En général, tracer la courbe représentant la somme des "distances" aux centres en les points non trimmés, en fonction du paramètre h , permet de déterminer le nombre de données aberrantes présentes dans l'échantillon et de les éliminer.

Les résultats que nous obtenons dans cette thèse sortent du cadre original de la méthode des k -moyennes trimmées au sens où nous remplaçons la norme Euclidienne par une divergence de Bregman. Ce point de vue avait été adopté dans le cas non trimmé par Banerjee et al. ; voir [BMDG05].

Une partie de notre apport consiste à dériver des vitesses non-asymptotiques pour ces méthodes de partitionnement trimmé. Nous utilisons à ces fins des outils de déviation de bornes pour les processus empiriques, faisant intervenir des entropies métriques par exemple. Ces résultats précisent les résultats asymptotiques obtenus dans le papier original de Cuesta-Albertos et al.

Nous adaptons les résultats de Cuesta-Albertos et al. (dans [CAGM97]) à des divergences de Bregman. La généralisation des propriétés de type existence d'un codebook optimal ou convergence presque sûre des codebooks empiriques optimaux n'est pas triviale. En effet, la difficulté majeure consiste à prouver que les codebooks empiriques optimaux (pour le partitionnement Bregman) sont bornés avec grande probabilité.

Cette méthode de partitionnement trimmé avec une divergence de Bregman d_ϕ est adaptée à des données échantillonnées selon des mélanges de lois dont la densité se réécrit sous la forme $x \mapsto \exp(-d_\phi(x, c) + C(x))$ où c est l'espérance de la loi. C'est le cas des lois Gaussiennes de même matrice de covariance, des lois de Poisson, binomiales, Gamma etc. Notre méthode appliquée à de telles données est illustrée dans la figure 1.11.

Lorsque les données sont échantillonnées dans le voisinage d'un compact (par exemple une sous-variété) selon une distribution P , le choix de k petit n'a plus de sens. On prend le point de vue de la quantification (c'est-à-dire k grand, mais petit par rapport à la taille de l'échantillon).

²⁷À chaque étape, on retire une proportion $1 - h$ des points, ceux qui sont les plus éloignés des centres selon une certaine "distance" (norme Euclidienne, divergence de Bregman etc.). Ensuite, les centres sont actualisés puis les points réintroduits avant de passer à l'étape suivante ; ce jusqu'à convergence de l'algorithme.




FIGURE 1.7 : Mélange Gaussien




FIGURE 1.8 : Mélange de Poisson




FIGURE 1.9 : Mélange binomial




FIGURE 1.10 : Mélange de Gamma



FIGURE 1.11 : Partitionnement trimmé avec une divergence de Bregman

En présence de données aberrantes, la méthode des k -moyennes trimmées échoue ; voir la figure 1.12.

La nouvelle approche que nous proposons consiste à faire de la quantification Bregman comme Banerjee et al. dans [BMDG05], mais avec une divergence de Bregman dépendante de la distribution P . La divergence de Bregman en question est associée à la fonction convexe $\psi_{P,h} : x \mapsto \|x\|^2 - d_{P,h,2}^2(x)$. Cette approche revient à minimiser en $\mathbf{t} = (t_1, t_2, \dots, t_k)$ le critère $P \min_{i \in [1,k]} \|u - m(t_i)\|^2 + v(t_i)$ où $m(t_i)$ et $v(t_i)$ correspondent à la moyenne et la variance de P restreintes à la boule centrée en t_i , de masse h pour P . Cette nouvelle méthode est une adaptation de la méthode des k -moyennes pour laquelle les centres sont à chercher parmi les moyennes locales. Le terme de variance locale a pour effet de rapprocher les centres des zones de forte densité.

Lorsque l'on implémente la version trimmée de cette méthode de quantification Bregman, on obtient la Figure 1.13.

La distance à la mesure introduite par Chazal et al. dans [CCSM11] est à la base de toutes les méthodes introduites et étudiées dans cette thèse. Elle possède de nombreux avantages. D'une part sa dépendance en un paramètre de masse h fait d'elle un objet géométrique robuste : Deux mesures proches au sens des métriques de Wasserstein ont des DTM proches. Ainsi, l'impact des données aberrantes dans un échantillon est minime pour le calcul de la DTM empirique. Dans ce cas, ce calcul se réduit à une recherche de q plus proches voisins, ce qui est peu coûteux. Aussi, ses diverses expressions permettent de faire le lien avec le critère des k -moyennes, des

FIGURE 1.12 : Quantification trimmée avec la méthode des k -moyennes trimméesFIGURE 1.13 : Quantification trimmée avec la divergence de Bregman associée à $\psi_{P,h}$

k -moyennes trimmées et même avec la log-vraisemblance trimmée.

Une expression de la DTM dans [CCSM11] est donnée en un point x de \mathbb{R}^d par $d_{P,h}^2(x) = P_{x,h} \|x - u\|^2$. La distribution $P_{x,h}$ correspond à la restriction de P à la boule Euclidienne de centre x et de masse h pour P . D'après Cuesta-Albertos et al., un minimiseur d'une telle fonction est en réalité un codebook optimal pour la méthode des k -moyennes trimmées, lorsque k vaut 1.

La perte des k -moyennes trimmées généralise la distance à la mesure à des codebooks \mathbf{c} de taille k , par la fonction $\mathbf{c} \mapsto P_{\mathbf{c},h} \min_{i \in \llbracket 1, k \rrbracket} \|c_i - u\|^2$. $P_{\mathbf{c},h}$ est la restriction de P à une union de k boules Euclidiennes centrées en les c_i , de masse totale h pour P . En remplaçant $\|c_i - u\|^2$ par $d_\phi(u, c_i)$, pour la divergence de Bregman d_ϕ , et les boules Euclidiennes par des boules de Bregman, on obtient une nouvelle généralisation de la distance à la mesure. Le codebook minimisant cette fonction correspond exactement au codebook optimal pour la méthode de partitionnement Bregman trimmé, c'est-à-dire pour notre méthode illustrée par la figure 1.11.

Nous mettons également à profit la convexité de la fonction $\psi_{P,h} : x \mapsto \|x\|^2 - d_{P,h,2}^2(x)$, construite à partir de la distance à la mesure. La méthode illustrée par la figure 1.13 repose en effet sur la méthode de quantification Bregman avec la divergence de Bregman associée à $\psi_{P,h}$. Il se trouve que la recherche d'un codebook optimal \mathbf{c}^* pour le critère correspondant revient à chercher la fonction puissance qui approche le carré de la distance à la mesure par au-dessus, au mieux (pour la norme $L_1(P)$). On appelle une telle approximation de la distance à la mesure k -PDTM (fonction k -puissance distance à la mesure). La méthode de quantification que nous proposons revient ainsi à déterminer un coresets²⁸ pour la distance à la mesure.

Dans le cas discret, en choisissant pour k la taille de l'échantillon n , la k -PDTM coïncide avec la q -witnessed distance de Guibas et al. dans [GMM11].

L'intérêt de notre approche est de fournir une approximation de la distance à la mesure dont les sous-niveaux sont des unions de k boules, avec k petit par rapport à n . On montre en particulier que le choix de k de l'ordre de $n^{\frac{d'}{d'+4}}$ est optimal, lorsque \mathcal{K} est une sous-variété de dimension d' . Les outils topologiques deviennent implémentables. La k -PDTM permet d'inférer la topologie d'un compact à partir d'un échantillon bruité de grand volume.

En théorie, on étudie la version non trimmée de la méthode de quantification Bregman

²⁸Le coresets d'une fonction construite sur un nuage de points est une fonction du même type, mais construite sur un nombre réduit de points k , et qui en est une bonne approximation.

associée à $\psi_{P,h}$. En pratique, on peut implémenter la version trimmée de la méthode de quantification Bregman. Les performances sont encore meilleures en présence de données aberrantes, puisque davantage de centres sont conservés.

Enfin, nous établissons un nouveau lien entre la distance à la mesure et des outils statistiques connus. Soit $\{P_\theta \mid \theta \in \Theta\}$, une famille de distributions paramétrée par la moyenne θ . On note p_θ la densité de P_θ . Soit P une mesure de probabilité sur Θ . On note $P_{\theta,h}$ la restriction de P au sous-niveau de $x \mapsto -\log(p_\theta(x))$, de masse h pour P . Lorsque le modèle considéré est une famille Gaussienne isotrope, ce sous-niveau est une boule de centre θ de masse h pour P .

La fonction définie en θ par $P_{\theta,h} - \log(p_\theta(u))$ est une généralisation de la distance à la mesure. Les deux notions coïncident lorsque l'on considère une famille Gaussienne isotrope. Dans le cas discret, un minimiseur θ de ce critère est en fait un maximum de la log-vraisemblance trimmée. On retrouve le maximum de vraisemblance lorsque $h = 1$. Cette généralisation peut-être vue comme une version continue de log-vraisemblance trimmée.

La méthode illustrée par la figure 1.15 consistait à approcher le graphe du carré de la DTM par l'enveloppe inférieure de k paraboles. On adapte cette méthode à notre généralisation de la log-vraisemblance trimmée, pour des modèles Gaussiens hétéroscédastiques. Après trimmage, on obtient la figure suivante 1.11. Cette méthode améliore l'algorithme tclust de [GEGMMI08] qui ne fonctionne plus lorsque les composantes se croisent, si l'on en croit la figure 1.14.




FIGURE 1.14 : Algorithme tclust de [GEGMMI08]




FIGURE 1.15 : Notre algorithme

FIGURE 1.16 : Partitionnement pour un modèle de mélange Gaussien hétéroscédastique.

Les contributions de cette thèse sont listées plus précisément dans la section suivante.

1.3 Présentation détaillée des chapitres

Comparaison de distributions au sens de Gromov-Wasserstein

Dans cette thèse, la première question que l'on traite est la suivante :

Deux nuages de points sont-ils issus de deux espaces métriques mesurés isomorphes ?

- Nous introduisons une famille d'ensembles de signatures $(\mathcal{S}_h)_{h \in [0,1]}$, indexée par un paramètre de masse h . Ce sont les *signatures-DTM*, elles sont construites à partir des fonctions distance à la mesure $d_{P,h,1}$.
 - Les éléments de \mathcal{S}_h sont des distributions sur \mathbb{R}^+ . Nous les comparons à l'aide de la distance de Wasserstein W_1 . Naturellement, nous en déduisons une famille de pseudo-

distances $(d_{\mathcal{S},h})_{h \in [0,1]}$ sur l'ensemble des espaces métriques mesurés quotienté par la relation d'isomorphisme.

- Pour tout h , nous montrons que \mathcal{S}_h satisfait les trois propriétés : *stabilité* (par rapport à la distance de Gromov-Wasserstein GW_1), *discrimination* (sous certaines alternatives) et *facilité d'implémentation* pour des espaces métriques discrets.
- Nous construisons une famille de *tests statistiques* $(\phi_{h,\rho})_{h \in [0,1], \rho}$ permettant de décider si deux nuages de points sont issus de deux espaces métriques mesurés isomorphes. Cette famille est indexée par un paramètre de masse h et un paramètre de sous-échantillonnage ρ .
 - La distribution de la statistique de test $T_{h,\rho}$ associée à $\phi_{h,\rho}$ est construite à partir de la pseudo-distance $d_{\mathcal{S},h}$. Sous l'hypothèse d'isomorphisme H_0 , nous étudions la distribution de $T_{h,\rho}$. En particulier, nous étudions sa *distribution limite*, et à taille d'échantillon fixé, son approximation avec une méthode de type *sous-échantillonnage*.
 - Le *niveau* des tests est étudié d'un point de vue asymptotique et leur *puissance* d'un point de vue non-asymptotique.
 - Nous fournissons un algorithme permettant d'implémenter ces tests statistiques et propose des heuristiques pour sélectionner en pratique les paramètres h et ρ .
 - La procédure que l'on propose est implémentée sur des exemples simulés et est comparée à d'autres alternatives éventuelles.

De façon annexe, nous démontrons quelques propriétés de discrimination des fonctions distance à la mesure. Nous donnons également une preuve détaillée du théorème de Gromov [GLP99, Théorème 3 1/2.5.] relatif à la reconstruction des espaces métriques mesurés à isomorphisme près, en se basant sur la preuve originale de Gromov.

Critères trimmés pour le partitionnement de données

Dans un second temps, nous nous focalisons sur la question suivante :

Comment partitionner au mieux un nuage de points en k groupes, en retirant éventuellement les données aberrantes ?

Les méthodes proposées sont particulièrement adaptées à certains modèles de mélange.

- Nous traitons d'abord le cas des mélanges de familles exponentielles pour lesquelles la densité se réécrit naturellement en fonction d'une divergence de Bregman. En ce sens, nous faisons le pont entre la méthode des *k -moyennes trimmées* (trimmed k -means) dans [CAGM97] et le partitionnement avec *divergence de Bregman* dans [BMDG05]. (Travaux réalisés en collaboration avec Aurélie Fischer et Clément Levrard)
 - La méthode de partitionnement étudiée repose sur la minimisation d'un critère de distorsion trimmé. On démontre l'existence d'un minimiseur c^* d'un tel critère lorsque la distribution P considérée a un moment d'ordre 1 fini.
 - Lorsqu'un n -échantillon de loi P est à disposition, et sous hypothèse d'unicité de c^* , nous étudions la *convergence d'un minimiseur du critère empirique* \hat{c}_n vers c^* . Aussi, sous des hypothèses de variance finie, nous dérivons des bornes sous-Gaussiennes pour l'*excès de distorsion* en \hat{c}_n .
 - Nous mettons en évidence un lien fort entre le critère de distorsion trimmé et la fonction distance à la mesure $d_{P,h,2}$.
 - Une adaptation de l'algorithme de Lloyd permet de calculer un minimiseur local du critère de distorsion empirique trimmé. Nous proposons une heuristique permettant de sélectionner le meilleur paramètre de trimmage en fonction des données. Finalement,

nous étudions les performances de ces méthodes sur des données échantillonnées selon des mélanges de lois de familles exponentielles, en présence de données aberrantes.

- Nous traitons ensuite le cas des mélanges de lois Gaussiennes hétéroscédastiques multidimensionnelles. Cette nouvelle technique repose sur la log-vraisemblance trimmée (trimmed log-likelihood).
 - La distance à la mesure $d_{P,h,2}$ peut être vue comme une version continue de la log-vraisemblance trimmée dans le cadre des lois Gaussiennes isotropes. Plus généralement, nous introduisons une *version continue de la log-vraisemblance trimmée*, comme généralisation de la DTM.
 - Nous développons une nouvelle méthode de partitionnement trimmé adaptée aux mélanges Gaussiens hétéroscédastiques. Cette méthode repose sur l'approximation de la version continue d'une certaine log-vraisemblance trimmée par l'enveloppe supérieure de k fonctions linéaires. Ces fonctions linéaires sont chacune paramétrées par un centre μ_i et une matrice de covariance Σ_i . Nous en déduisons un critère de coût fonction de k centres μ_i et k matrices de covariance Σ_i .
 - Lorsque $k = 1$ et $P = \mathcal{N}(\mu_0, \Sigma_0)$, nous établissons l'expression des minimiseurs du critère de coût en fonction de μ_0 et Σ_0 .
 - Nous fournissons un algorithme permettant de calculer un minimum local du critère de coût, ce qui fournit un partitionnement des données. Nous comparons ensuite ses performances à celles obtenues par l'algorithme tclust de García-Escudero et al. dans [GEGMMI08].

De façon annexe, nous revenons sur les méthodes dites de chaînage visant à fournir des bornes de déviation pour les processus empiriques. Parmi les outils utilisés, on peut citer l'entropie métrique avec la dimension de Vapnik-Chervonenkis, l'intégrale de Dudley, le principe de symétrisation avec les moyennes de Rademache etc. Nous y rappelons les méthodes développées dans [BBL05, Section 3] et [BLM13] par exemple. Nous les illustrons par divers exemples utilisés dans cette thèse. La plupart de ces exemples ont été obtenus par Clément Levard, dans le cadre de nos collaborations.

Une nouvelle méthode de quantification basée sur la distance à la mesure

Dans un dernier temps, nous apportons une réponse à la question suivante :

Étant donné un nuage de points tirés dans un voisinage d'un compact \mathcal{K} , avec éventuellement quelques données aberrantes, comment choisir k points pour approcher au mieux la distance à \mathcal{K} ?

Cette question est liée au problème de quantification d'une distribution P , avec comme objectif supplémentaire la fonction distance à son support \mathcal{K} . (Travaux réalisés en collaboration avec Clément Levard)

- Nous étudions la méthode de partitionnement de [BMDG05] pour la divergence de Bregman associée à la fonction convexe $x \mapsto \|x\|^2 - d_{P,h,2}^2(x)$. Cela revient à minimiser en $\mathbf{t} = (t_1, t_2, \dots, t_k)$ le critère $P \min_{i \in [1,k]} \|u - m(t_i)\|^2 + v(t_i)$ où $m(t_i)$ et $v(t_i)$ correspondent à la moyenne et la variance de P restreintes à la boule centrée en t_i , de P -masse h . Nous adoptons le point de vue de la quantification (i.e. k est grand).
 - La minimisation du critère de coût revient au problème suivant : chercher la meilleure approximation du carré de la DTM $d_{P,h,2}^2$ par une union de k paraboles, par au dessus,

- pour la norme $L_1(P)$. On appelle k -PDTM une telle approximation de la DTM. Ses sous-niveaux sont des unions de k boules.
- On se place dans le contexte où seule une modification Q de la distribution P (dont le support \mathcal{K} est à inférer) est à disposition. Nous étudions l'approximation du carré de la DTM associée à P par le carré de la k -PDTM associée à Q , dans $L_1(P)$.
 - Nous calculons des bornes pour la norme infinie de la différence entre la k -PDTM associée à Q et la distance au compact \mathcal{K} . En particulier, l'information topologique contenue dans \mathcal{K} peut être reconstruite à partir de la k -PDTM associée à Q , pourvu que h soit suffisamment petit et Q proche de P au sens de Wasserstein.
 - La k -PDTM empirique est définie comme la k -PDTM associée à Q_n , la distribution empirique associée à un n -échantillon de loi Q . Nous étudions le cas particulier où la distribution Q est obtenue comme la convolution de P par une mesure sous-Gaussienne de variance σ^2 .
 - La vitesse d'approximation de la k -PDTM associée à Q par la k -PDTM empirique dans $L_1(P)$ est de l'ordre de $\frac{\sqrt{kd}}{h\sqrt{n}} + \frac{\sigma}{\sqrt{h}}$. Le premier terme correspond aux vitesses lentes usuelles pour la méthode des k -moyennes. Le terme supplémentaire est un terme de biais du fait que l'approximation soit mesurée dans l'espace $L_2(P)$ plutôt que dans $L_2(Q)$.
 - Une optimisation en k des diverses bornes obtenues incite à choisir k de l'ordre de $n^{\frac{d'}{d'+4}}$, lorsque \mathcal{K} est une sous-variété de dimension d' .
 - L'algorithme de Banerjee [BMDG05] pour la divergence de Bregman associée à la fonction convexe $x \mapsto \|x\|^2 - d_{Q_n, h, 2}^2(x)$ fournit un minimum local pour le critère empirique. Nous implémentons également la version Bregman trimmée exposée plus haut. Dans les deux cas, une heuristique permet de sélectionner le nombre de points à trimmer, i.e. ceux qui sont considérés comme des données aberrantes.

De façon annexe, nous étudions la convexité de l'ensemble des moyennes locales associées à une distribution P , pour un paramètre de masse h fixé. Par moyenne locale, on entend espérance d'une sous-mesure de P de masse h . La convexité est assurée en particulier lorsque P est une mesure Borélienne de \mathbb{R}^d qui est continue. Ces résultats permettent de faire le lien plus facilement entre les différentes définitions de la k -PDTM.

Pour aller plus loin :

Les tests statistiques d'isomorphisme introduits dans cette thèse reposent sur la distance à la mesure de paramètre 1. Ces tests sont indexés par un paramètre de masse h . La distance à la mesure jouit d'un certain nombre de propriétés qui nous permettent d'avoir des garanties sur son efficacité pour les tests que l'on développe. Une question naturelle est de se demander s'il existe d'autres fonctions qui vérifient ces propriétés, et qui seraient donc tout autant adaptées au test d'isomorphisme. Il est intéressant de disposer de différentes fonctions puisqu'il est possible que la distance à la mesure soit parfois incapable de distinguer que deux distributions sont non isomorphes, alors qu'une autre fonction serait capable de les distinguer. D'autres pistes d'amélioration sont à considérer. Dans certains cas, nos tests ne fonctionnent pas, en particulier lorsque les signature-DTM sont des mesures de Dirac. C'est le cas des distributions uniformes sur une sphère. Dans ce cas, une piste serait de convoler les deux distributions par une loi normale, on en déduirait un test statistique qui fonctionnerait même dans les cas dégénérés. Une autre piste d'amélioration du test consisterait à renormaliser la statistique de test par une

certainne quantité, de sorte à autoriser des dilatations entre les deux distributions à comparer.

D'autre part, on rappelle qu'à partir d'un même échantillon, plusieurs p -valeurs différentes peuvent dériver d'un même test statistique d'isomorphisme. Cela est dû au fait que la statistique de test est construite sur un sous-ensemble de l'échantillon. Les p -valeurs sont donc aléatoires. Les dépendances entre deux p -valeurs associées au même test et au même échantillon sont faibles pourvu qu'elles soient construites à partir de deux parties distinctes de l'échantillon. Il serait intéressant de tirer parti de cette propriété. Par exemple, le test de Kolmogorov-Smirnov appliqué à plusieurs p -valeurs pour le même h permet de fabriquer un test un peu plus puissant en pratique. D'autre part, une procédure de type Benjamini-Hochberg permet de sélectionner les paramètres h pour lesquels la signature-DTM discrimine le mieux les deux distributions. A priori, la procédure de Benjamini-Hochberg est valable pour des p -valeurs indépendantes. Dans notre cas, les p -valeurs ne sont pas indépendantes, mais leurs dépendances sont faibles. Il serait intéressant de développer des outils théoriques permettant de justifier ces méthodes expérimentales.

Les vitesses usuelles de convergence des codebooks empiriques vers les codebooks optimaux sont de l'ordre de $n^{-\frac{1}{2}}$ pour la méthode des k -moyennes. Nous montrons dans cette thèse que les mêmes vitesses sont obtenues pour la méthode de partitionnement trimmé avec une divergence de Bregman. Dans [Lev14], Levrard propose des vitesses plus rapides de l'ordre de n^{-1} , sous des hypothèses de conditions de marges. Il serait intéressant de chercher à adapter ces vitesses à notre cadre plus général. La théorie des méthodes de partitionnement trimmées étudiées dans cette thèse repose sur des hypothèses de sous-gaussianité. Pourtant, expérimentalement ces méthodes s'adaptent à des distributions à queues lourdes, on pense aux lois de Cauchy. Nous pourrions essayer d'adapter les résultats de cette thèse aux distributions à queues lourdes.

Nous proposons dans cette thèse une méthode permettant de partitionner des données échantillonnées selon un modèle Gaussien hétéroscédastique, en présence d'éventuelles données aberrantes. Une telle méthode pourrait être utilisée pour inférer les moyennes et matrices de covariance des distributions du mélange. Cela semble possible lorsque les moyennes sont dans des zones de faibles densité pour les autres distributions du mélange, que h est suffisamment petit, et que la taille de l'échantillon suffisamment grande. Il serait intéressant de donner une réponse probabiliste ou statistique à cette question.

Approximer un nuage de points par un ensemble de k -points peut être fait au sens de la distance de Hausdorff, on obtient alors un *réseau*. Lorsque cette approximation est faite par une méthode de type k -moyennes, la distance qui est minimisée est la distance de Wasserstein. Notre méthode approxime elle aussi la distribution empirique associée à un nuage de points, tout en étant plus robuste aux données aberrantes. Ces approximations peuvent permettre en utilisant les méthodes d'inférence topologiques de reconstruire les sous-variétés sous-jacentes.

Une adaptation simple de notre procédure de recherche de codebooks optimaux pour la k -PDTM pourrait permettre de rechercher les modes d'une distribution, en faisant décroître le paramètre h dans l'algorithme de recherche du codebook optimal. Il serait intéressant de faire l'étude théorique et pratique d'une telle méthode. Enfin, une adaptation de la méthode de single-linkage avec la k -PDTM permettrait de fournir une méthode de classification hiérarchique adaptée à la présence de données aberrantes.

La distance à la mesure est une version continue du critère des moindres carrés trimmés. De façon plus globale, des généralisations de la distance à la mesure, construites à partir de sous-mesures, pourraient être utilisées pour fournir des versions continues d'objets discrets trimmés. Ceci permettrait d'en faciliter l'étude statistique, comme dans cette thèse.

Similarity of distributions according to the Gromov-Wasserstein distance

Dans ce chapitre, nous traitons de la comparaison de distributions au sens de Gromov-Wasserstein. En ce sens, nous développons des outils et méthodes permettant de décider si deux espaces métriques mesurés sont isomorphes. Dans un premier temps, nous introduisons une famille de signatures construites à partir de la famille des fonctions distance à la mesure $(d_{P,h,1})_{h \in [0,1]}$, les *signatures-DTM*. Nous étudions leurs propriétés de stabilité et de discrimination. Dans un second temps, nous développons une famille de *tests statistiques* permettant de décider si deux échantillons de points sont issus de deux espaces métriques mesurés isomorphes. Ces tests reposent sur des pseudo-distances dérivées des signatures-DTM. Ce travail a fait l'objet d'un papier [Bré18], soumis à *Electronic Journal of Statistics*. Une version géométrique du papier est disponible dans [Bré16].

Contents

2.1	Discriminative tools for comparing probability distributions	29
2.1.1	From comparison of compact sets to comparison of probability distributions	29
2.1.2	Equality of probability distributions up to rigid transformations	35
2.2	Statistical test of similarity in the Gromov-Wasserstein sense	40
2.2.1	A test of asymptotic level α	42
2.2.2	The power of the test	45
2.2.3	The algorithm	46
2.2.4	A discussion of other methods	49
2.2.5	On the selection of the parameters	52
2.3	Proofs	53
2.3.1	Proofs for Section 2.1	53
2.3.2	Proofs for Section 2.2	61

In this chapter, we deal with the comparison of distributions in the Gromov-Wasserstein sense. To this aim, we provide tools and methods to decide whether two metric-measure spaces are isomorphic or not. Firstly, we introduce a family of signatures built from the family of the distance-to-measure functions $(d_{P,h,1})_{h \in [0,1]}$, the *DTM-signatures*. We consider their stability and discriminative properties. Secondly, we derive a family of statistical tests that amount to decide whether two samples of points have been generated from two isomorphic metric-measure spaces. Such statistical tests are based on pseudo-metrics derived from the DTM-signatures.

2.1 Discriminative tools for comparing probability distributions

In this section, we expose and develop tools to compare probability measures. Over a first phase, we consider two distributions defined on the same Borel σ -algebra. The pseudo-metrics and metrics we exhibit aim at measuring proximity of two such distributions, or proximity of their supports. Over a second phase, we consider two metric spaces equipped with a distribution (i.e. two metric-measure spaces). The notion of proximity is then measured in terms of the Gromov-Wasserstein metric, developed by Mémoli in [Mém11] (see Definition 2.10). According to this metric, two metric-measure spaces are considered the same when they are equal up to an isomorphism, i.e. a rigid transformation, as defined in Definition 2.9.

2.1.1 From comparison of compact sets to comparison of probability distributions

Set P , a probability distribution defined on the Borel σ -algebra of a metric space (\mathcal{X}, δ) . The support of P is defined as the largest closed subset of \mathcal{X} with P -mass¹ 1. The support of P is denoted by $\text{Supp}(P)$.

A first strategy to compare two probability distributions consists in comparing their supports. Whenever the supports are compact sets, their proximity is measured in terms of the Hausdorff distance.

Distance to a compact set and Hausdorff distance

Set \mathcal{K} , a compact subset of (\mathcal{X}, δ) . The *distance to the compact set \mathcal{K}* , is a function $d_{\mathcal{K}} : \mathcal{X} \rightarrow \mathbb{R}_+$ defined for every $x \in \mathcal{X}$ by $d_{\mathcal{K}}(x) = \inf \{\delta(x, y) \mid y \in \mathcal{K}\}$. For every $\epsilon \geq 0$, the ϵ -offset of \mathcal{K} is defined as the ϵ -sublevel set of $d_{\mathcal{K}}$ by

$$\mathcal{K}^\epsilon = \{x \in \mathcal{X} \mid d_{\mathcal{K}}(x) \leq \epsilon\} = \bigcup_{x \in \mathcal{K}} \overline{\mathcal{B}}(x, \epsilon),$$

where $\overline{\mathcal{B}}(x, \epsilon) = \{y \in \mathcal{X} \mid \delta(x, y) \leq \epsilon\}$ denotes the closed ball centered at x with radius ϵ in the space (\mathcal{X}, δ) .

A compact set \mathcal{K} can be recovered from $d_{\mathcal{K}}$ since it coincides with its 0-offset:

$$\mathcal{K} = d_{\mathcal{K}}^{-1}(0) = \{x \in \mathcal{X} \mid d_{\mathcal{K}}(x) = 0\}.$$

As a consequence, two compact sets \mathcal{K} and \mathcal{K}' are equal if and only if $d_{\mathcal{K}}$ and $d_{\mathcal{K}'}$ are the same. The Hausdorff distance d_H is defined after such a consideration by

$$d_H(\mathcal{K}, \mathcal{K}') = \sup_{x \in \mathcal{X}} |d_{\mathcal{K}} - d_{\mathcal{K}'}|(x).$$

Equivalently, a definition for the Hausdorff distance can be given in terms of offsets:

$$d_H(\mathcal{K}, \mathcal{K}') = \inf \{\epsilon > 0 \mid \mathcal{K} \subset \mathcal{K}'^\epsilon \text{ and } \mathcal{K}' \subset \mathcal{K}^\epsilon\}.$$

See e.g. [Aam17] for the equivalence of the definitions and for additional properties satisfied by the Hausdorff distance. In a nutshell, two compact sets \mathcal{K} and \mathcal{K}' are equal if and only if their Hausdorff distance is zero, i.e. if and only if $d_{\mathcal{K}}$ and $d_{\mathcal{K}'}$ coincide everywhere. Moreover, \mathcal{K} coincides with the set of points $x \in \mathcal{X}$ for which $d_{\mathcal{K}}$ is zero.

The Hausdorff distance is perfectly adapted to topological data analysis. Indeed, if $D(\mathcal{K})$ denotes the persistence diagram associated with $(\mathcal{K}^\epsilon)_{\epsilon \geq 0}$, the family of sublevel sets of $d_{\mathcal{K}}$, and dB is the bottleneck distance between persistence diagrams, we have:

$$\text{dB}(D(\mathcal{K}), D(\mathcal{K}')) \leq d_H(\mathcal{K}, \mathcal{K}'). \quad (2.1)$$

It comes as a direct application of the stability theorem for persistence diagrams (Theorem 1.4), that are defined in Section 1.1.3.

Now, we take the point of view of topological inference. In that, we consider a set of n points \mathcal{K}_n in \mathcal{X} and intend to recover topological information for \mathcal{K} (number of connected components, holes, voids...) from \mathcal{K}_n . If $d_H(\mathcal{K}_n, \mathcal{K})$ is small, then according to (2.1), the topological information contained in $D(\mathcal{K}_n)$ will roughly be the same as the information contained in $D(\mathcal{K})$.

¹We say that a Borel set A has P -mass $h \in [0, 1]$ whenever $P(A) = h$.




Figure 2.1: Distance to \mathcal{K}




Figure 2.2: Distance to \mathcal{K}_n




Figure 2.3: DTM to Q_n






Figure 2.4: Barcodes associated to Figures 2.1, 2.2 and 2.3

Nonetheless, if some points in \mathcal{K}_n are far from \mathcal{K}^2 , this assertion turns out to be wrong. Indeed, in this case, the distance functions $d_{\mathcal{K}}$ and $d_{\mathcal{K}_n}$ are very different, as highlighted by Figures 2.1 and 2.2. As a consequence, the barcodes (that contain the same information as the persistence diagrams) are completely different; see Figure 2.4. Then, the topology of \mathcal{K} cannot be recovered from $d_{\mathcal{K}_n}$.

An alternative point of view consists in replacing compact sets by distributions. For P a probability measure with support \mathcal{K} , we intend to recover the topology of \mathcal{K} from a distribution Q_n close to P . As an example, we consider the distribution $P = \frac{1}{4} (\delta_0 + \delta_{0.9} + \delta_1 + \delta_{1.1})$, with δ_x the *Dirac measure*³ centered on x . The measure P is supported on the set $\mathcal{K} = \{0, 0.9, 1, 1.1\}$ represented in Figure 2.1.

If Figure 2.2, we observe a set \mathcal{K}_n . Then we can set $Q_n = \frac{1}{n} \sum_{x \in \mathcal{K}_n} \delta_x$. Some elements in \mathcal{K}_n are far from \mathcal{K} . Nonetheless, the distribution Q_n remains close to P in terms of the Wasserstein metric, as defined by Definition 1.2. As a consequence, as highlighted by Figure 2.3, we see that $d_{\mathcal{K}}$ is well approximated by the distance-to-measure (DTM) to Q_n . In addition, the barcode associated with \mathcal{K} is close to the barcode associated with \mathcal{K}_n , see Figure 2.4.

Though, in presence of outliers, a strategy to compare compact sets or measures consists in replacing the function distance to a compact set with the function distance-to-measure, as defined below.

Distance-to-measure and Wasserstein metrics

The distance-to-measure, introduced by Chazal et al. in [CCSM11] is a generalization of the function distance to a compact set that depends on a mass parameter $h \in [0, 1]$ and is defined as follows.

²Such points are called *outliers*.

³Set (Ω, \mathcal{F}) a measurable space, the Dirac measure centered on $x \in \Omega$ is a probability measure defined for all $A \in \mathcal{F}$ by $\delta_x(A) = 1$ if $x \in A$ and $\delta_x(A) = 0$ if $x \notin A$.

Definition 2.1 (Chazal et al. 2011, see [CCSM11]). Set P , a Borel probability measure on a metric space (\mathcal{X}, δ) and h a mass parameter in $[0, 1]$.

The *pseudo-distance* with mass parameter $l \in (0, 1)$ is the function $\delta_{P,l}: \mathcal{X} \rightarrow \mathbb{R}_+$ defined for every $x \in \mathcal{X}$ by

$$\delta_{P,l}(x) = \inf\{\epsilon > 0 \mid P(\overline{\mathcal{B}}(x, \epsilon)) > l\}.$$

Then, the *distance-to-measure* (DTM) associated to P with mass parameter h is the function $d_{P,h}: \mathcal{X} \rightarrow \mathbb{R}_+$ ⁴ defined for every $x \in \mathcal{X}$ by

$$d_{P,h}(x) = \frac{1}{h} \int_{l=0}^h \delta_{P,l}(x) dl$$

More generally, the DTM with parameter $r \geq 1$ is defined for every $x \in \mathcal{X}$ by

$$d_{P,h,r}(x) = \left(\frac{1}{h} \int_{l=0}^h \delta_{P,l}^r(x) dl \right)^{\frac{1}{r}}$$

In this chapter, when r is not specified for the DTM, the parameter equals to $r = 1$.

The DTM to a distribution P approximates the function distance to its support, $\text{Supp}(P)$, in the sense that for every $x \in \mathcal{X}$:

$$\lim_{h \rightarrow 0} d_{P,h}(x) = d_{\text{Supp}(P)}(x).$$

In [CCSM11, Corollary 4.8 and Proposition 4.9], upper-bounds are derived for the difference $\|d_{P,h} - d_{\text{Supp}(P)}\|_\infty$, in the context of uniform distributions on submanifolds. The broader framework of (a, b) -standard distributions⁵, for some a and b in \mathbb{R}_+ , is also considered.

By analogy with the Hausdorff distance construction, pseudo-metrics between two measures P and Q can be derived from the corresponding DTM. A natural definition for such pseudo-metrics is given by the maximal difference between the two DTM, $\|d_{P,h,r} - d_{Q,h,r}\|_\infty$. This family of pseudo-metrics are actually to relate with the family of Wasserstein distances $(W_r)_{r \geq 1}$.

Proposition 2.2 (Stability, in [CCSM11] for \mathbb{R}^d , in [Buc14] for metric spaces). *If P and Q are two probability distributions on the space (\mathcal{X}, δ) , it comes that*

$$\|d_{P,h,r} - d_{Q,h,r}\|_\infty \leq \frac{1}{h^{\frac{1}{r}}} W_r(P, Q),$$

where the supremum is taken on the whole space \mathcal{X} .

We recall that the Wasserstein metric W_r with parameter $r \geq 1$ is defined for two probability distributions P and Q with a finite r moment⁶ by:

$$W_r^r(P, Q) = \inf_{\pi \in \Pi(P, Q)} \mathbb{E}_{(X, Y) \sim \pi} [\delta^r(X, Y)].$$

The set of *transport plans* $\Pi(P, Q)$ corresponds to the set of all distributions on $\mathcal{X} \times \mathcal{X}$ of a random vector (X, Y) , for X generated according to P and Y according to Q .

⁴When $h = 1$, the DTM might be equal to infinity. When $h \in [0, 1)$, the DTM is always finite and takes values in \mathbb{R}_+ .

⁵A distribution P is (a, b) -standard if $P(\overline{\mathcal{B}}(x, r)) \geq \min\{1, ar^b\}$ for every $x \in \text{Supp}(P)$ and $r \geq 0$.

⁶ P has a finite r moment if for some $x_0 \in \mathcal{X}$, $P\delta(x_0, u)^r < \infty$.

Set $\mathbb{X}_n = \{X_1, X_2, \dots, X_n\}$ an n -sample from P and $P_n = \frac{1}{n} \sum_{i=1}^n \delta_{X_i}$, the corresponding empirical distribution. According to [BL14, Theorem 2.13], for every $r \geq 1$ $W_r(P_n, P)$ converges to 0 a.e., provided that P has a finite moment of order r .

As a consequence, it is relevant to approximate the DTM from \mathbb{X}_n via its empirical counterpart, the *empirical distance-to-measure*, that is obtained by replacing P with P_n in the definition of the DTM. In that sense, we say that the distance-to-measure is stable under sampling.

Moreover, it is possible to obtain upper-bounds for the variation $\|d_{P,h,r} - d_{Q,h,r}\|_\infty$ when Q is a modification of P given as the convolution of P with a distribution \mathcal{N} with finite r moment, by $Q = P * \mathcal{N}$. That is, Q is the distribution of the random variable $Y = X + Z$ where X is generated according to P and Z according to \mathcal{N} . Then, it comes that

$$W_r^r(P, Q) \leq \mathbb{E} [\|Z\|^r] = \mathcal{N}\|u\|^r.$$

A sample from Q consists in data points generated from P with a slight error of measurement given by the distribution \mathcal{N} . Often, \mathcal{N} is a Gaussian distribution.

Another type of transformations consists in replacing the measure P with the distribution $Q = \alpha P + (1 - \alpha)\mathcal{C}$ for α close to 1 and \mathcal{C} a distribution with bounded support, in a ball $\mathcal{B}(0, R)$. For such a distribution Q , it comes that

$$W_r^r(P, Q) \leq (1 - \alpha)R^r.$$

A sample from Q is composed of many points generated according to P and a few points generated from \mathcal{C} . When the distribution \mathcal{C} is uniform on an open set, we use the term “clutter noise”.

Previously, we mentioned two very important properties satisfied by the distance to a compact set. Firstly, two compact sets \mathcal{K} and \mathcal{K}' are equal if and only if $d_{\mathcal{K}} = d_{\mathcal{K}'}$. Secondly, it is possible to recover the compact set \mathcal{K} from the function distance to the compact set, as $d_{\mathcal{K}}^{-1}(0) = \mathcal{K}$. We will prove that similar properties are satisfied by the DTM.

The question of the determination of a measure from its DTM is of interest. Some work has been done in this direction for discrete measures; see [Buc14, Theorem 3.21]. Firstly, we prove that for some particular spaces, the knowledge of the DTM at every point x and for every $h \in (0, 1)$ determines the distribution. Secondly, we prove that uniform distributions with regular⁷ support can be recovered from the DTM associated with a mass parameter $h > 0$ small enough.

Proposition 2.3. *We define maps ϕ and ψ for every probability measure P such that $P\|u\| < \infty$ by:*

$$\phi(P) = (d_{P,h}(x))_{h \in (0,1), x \in \mathcal{X}}$$

and

$$\psi(P) = (P(\overline{\mathcal{B}}(x, r)))_{r \in \mathbb{R}_+, x \in \mathcal{X}}.$$

Then, the map ϕ is injective if and only if the map ψ is injective.

The proof of Proposition 2.3 is available in Section 2.3.1.

Concretely, it means that spaces (\mathcal{X}, δ) equipped with the Borel σ -algebra $\sigma(\mathcal{X})$, that satisfy the property:

⁷Here a set is *regular* if it has a positive reach, as defined by Equation (2.2).

– Every distribution on $\sigma(\mathcal{X})$ is determined by its values on balls.
 coincide with the spaces that satisfy:

– Every distribution on $\sigma(\mathcal{X})$ is determined by its DTM for all $h \in (0, 1)$ and $x \in \mathcal{X}$.

As noticed in [BL16], Euclidean spaces \mathbb{R}^d satisfy these assertions, but there exists metric spaces for which these assertions are not satisfied.

In the following, we will establish a stronger identifiability result in the following specific framework.

Set O , a non-empty bounded open subset of \mathbb{R}^d satisfying $(\overline{O})^\circ = O$ ⁸. Then, the *uniform measure* P_O is defined for every Borel set A of \mathbb{R}^d , by:

$$P_O(A) = \frac{\text{Leb}_d(O \cap A)}{\text{Leb}_d(O)},$$

with Leb_d the Lebesgue measure on \mathbb{R}^d .

The medial axis of O , $\mathcal{M}(O)$ is defined as the set of points in O having at least two projections onto its boundary ∂O . That is,

$$\mathcal{M}(O) = \{y \in O \mid \exists x', x'' \in \partial O, x' \neq x'', \|y - x'\|_2 = \|y - x''\|_2 = d_{\partial O}(y)\},$$

with $d_{\partial O}(y) = \inf\{\|x - y\|_2 \mid x \in \partial O\}$.

The *reach* of O , $\text{Reach}(O)$ is defined as the distance between its boundary ∂O and its medial axis $\mathcal{M}(O)$. That is,

$$\text{Reach}(O) = \inf\{\|x - y\|_2 \mid x \in \partial O, y \in \mathcal{M}(O)\}. \tag{2.2}$$

If \mathcal{K} is a compact subset of \mathbb{R}^d , it is standard to define its reach as $\text{Reach}(\mathcal{K}^c)$, the reach of its complement in \mathbb{R}^d . See [Fed96] to get more familiar with these notions.




Figure 2.5: Medial Axis and Reach of an open set

The notion of medial axis is to relate with the notion of skeleton, as enhance by the following lemma from Lieutier.

Lemma 2.4 ([Lie04]). *The skeleton $Sk(O)$ of the open set O is defined as the set of centres of maximal balls (for the inclusion) included in O . Then, it comes that*

$$\mathcal{M}(O) \subset Sk(O) \subset \overline{\mathcal{M}(O)}.$$

⁸Recall that the notation \overline{A} stands for the closure of A , and A° for its interior.

In Section 2.3.1, we give a proof of Lemma 2.4. This proof was obtained in collaboration with Frédéric Chazal.

Note that when P is a finite set, these inclusions are actually equalities. Nonetheless, there exists sets for which these inclusions are not equalities. For instance, in [Lie04], Lieutier proves that the medial axis of an ellipse is given by a segment without its extremities.

The knowledge of the distance-to-measure at every $x \in \mathbb{R}^d$ for two distributions P_O and $P_{O'}$ is discriminative, provided that h is small enough.

Proposition 2.5. *Let O and O' be two non-empty bounded open subsets of \mathbb{R}^d with positive reach, such that $O = (\overline{O})^\circ$ and $O' = (\overline{O'})^\circ$. Let $h > 0$ be such that*

$$h \leq \min \left(\text{Reach}(O)^d, \text{Reach}(O')^d \right) \frac{\omega_d}{\text{Leb}_d(O)}.$$

with $\omega_d = \text{Leb}_d(\mathcal{B}(0, 1))$, the Lebesgue volume of the unit d -dimensional ball. If for every $x \in \mathbb{R}^d$

$$d_{P_O, h}(x) = d_{P_{O'}, h}(x),$$

then $P_O = P_{O'}$.

The proof of Proposition 2.5 is available in Section 2.3.1 and is based on the two following propositions.

On the one hand, the set of minimisers of the DTM of P_O coincides with the set of points in O which distance to O is larger than some radius $\epsilon(h, O)$. This radius is defined for $h \in (0, 1)$ by

$$\epsilon(h, O) = \left(\frac{h \text{Leb}_d(O)}{\omega_d} \right)^{\frac{1}{d}}.$$

Such an $\epsilon(h, O)$ corresponds to the radius of a ball included in O , with P_O -mass h .

Proposition 2.6. *The constant $d_{\min} = \frac{d}{d+1} \left(\frac{h \text{Leb}_d(O)}{\omega_d} \right)^{\frac{1}{d}}$ is a lower bound for the DTM of P_O over \mathbb{R}^d . Moreover, the set of points attaining this bound is exactly $O_{\epsilon(h, O)}$. where*

$$O_\epsilon = \left\{ x \in O \mid \inf_{y \in \partial O} \|x - y\|_2 \geq \epsilon \right\}.$$

The proof of Proposition 2.6 is to be found in Section 2.3.1.

On the other hand, O can be recovered as the $\epsilon(h, O)$ -offset of the set of minimizers of the DTM of P_O .

Proposition 2.7. *If $\text{Reach}(O) \geq \epsilon(h, O)$, then:*

$$\{x \in \mathbb{R}^d \mid d_{P_O, h}(x) = d_{\min}\}^{\epsilon(h, O)} = O.$$

The proof of Proposition 2.7 is to be found in Section 2.3.1.

Many metrics, pseudo-metrics and dissimilarity measures have been studied in the literature to compare distributions defined on the same σ -algebra. Among others, we can cite the total-variation distance, the Kullback-Leibler divergence, the Prokhorov metric, the Kolmogorov distance and the maximum mean discrepancy.

In this section, we have proposed alternatives to these tools based on the DTM. The pseudo-metrics we consider are specially designed for the context of topological inference with data corrupted by outliers. The pseudo-metrics aim at detecting equality of two probability distributions.

2.1.2 Equality of probability distributions up to rigid transformations

As aforementioned, the comparison of two datasets may be compromised when the data sets have not been generated on the same space, or have not been stored within the same coordinate system. In the context of compact sets comparison, the Gromov-Hausdorff metric GH was introduced in 1975 by Edwards [AE75] and reintroduced in 1981 by Gromov [GLP81, Gro81]. Moreover, according to [CdSGO16] or [CGLM15, Equation (2.3)], the persistence diagrams associated with compact sets \mathcal{K} and \mathcal{K}' that are close in the Gromov-Hausdorff sense, roughly contain the same topological information, in that:

$$dB(D(\mathcal{K}), D(\mathcal{K}')) \leq GH(\mathcal{K}, \mathcal{K}').$$

When dealing with distributions, it is common to work with the following setting.

Definition 2.8. A *metric-measure space (mm-space)* is a triple (\mathcal{X}, δ, P) , with (\mathcal{X}, δ) a metric space and P a probability measure on \mathcal{X} equipped with its Borel σ -algebra.

Definition 2.9 (Isomorphism between mm-spaces). Two mm-spaces (\mathcal{X}, δ, P) and (\mathcal{Y}, γ, Q) are said to be *isomorphic* if:

- Their exist Borel sets $\mathcal{X}_0 \subset \mathcal{X}$ and $\mathcal{Y}_0 \subset \mathcal{Y}$ such that $P(\mathcal{X} \setminus \mathcal{X}_0) = 0$ and $Q(\mathcal{Y} \setminus \mathcal{Y}_0) = 0$.
- There is some one-to-one and onto isometry $\phi : \mathcal{X}_0 \rightarrow \mathcal{Y}_0$.
- The map ϕ preserves measures, i.e. $Q(\phi(A \cap \mathcal{X}_0)) = P(A \cap \mathcal{X}_0)$ for every Borel set A of \mathcal{X} .

Such a map ϕ is called an *isomorphism* between the mm-spaces (\mathcal{X}, δ, P) and (\mathcal{Y}, γ, Q) .

In this section, we address the question of the comparison of general mm-spaces, up to an isomorphism. In other terms, we aim at designing a metric or at least a pseudo-metric on the quotient space of mm-spaces by the relation of isomorphism. A suitable pseudo-distance should be stable under some perturbations, under sampling, discriminative and easy to implement when dealing with discrete spaces.

From metrics to pseudo-metrics to compare mm-spaces

A first characterisation of mm-spaces is given in [GLP99]. In its Theorem 3 $\frac{1}{2}$.5, Gromov proves that any mm-space can be recovered, up to an isomorphism, from the knowledge, for all size N , of the distribution of the $N \times N$ -matrix of distances associated to an N -sample. A proof of this theorem is available in the appendix, in Section 2.A. More recently, in [Mém11], Mémoli introduced metrics on the quotient space of mm-spaces by the relation of isomorphism, the Gromov–Wasserstein distances.

Definition 2.10 (Gromov–Wasserstein distance). The *Gromov–Wasserstein distance* between two mm-spaces (\mathcal{X}, δ, P) and (\mathcal{Y}, γ, Q) denoted $GW(P, Q)$ is defined by:

$$GW(P, Q) = \inf_{\pi \in \Pi(P, Q)} \frac{1}{2} \mathbb{E}_{(X', Y'), (X, Y) \sim \pi} [|\delta(X, X') - \gamma(Y, Y')|].$$

Again, $\Pi(P, Q)$ stands for the set of transport plans between P and Q , and (X', Y') and (X, Y) are independent random vectors with distribution π .

Unfortunately, even when dealing with discrete mm-spaces, the computation of these Gromov–Wasserstein distances is extremely costly. An alternative is to assign a *signature* to every mm-space. A signature is an element (in \mathbb{R} , \mathbb{R}^d , or a distribution on \mathbb{R}) assigned to a mm-space such that the same element is assigned to isomorphic mm-spaces. The mm-spaces are then compared through their signatures. In [Mém11], Mémoli gives an overview of such signatures, as for instance shape distribution, eccentricity or what he calls local distribution of distances.

In this section, we introduce a new signature that is a probability measure on \mathbb{R} , and we propose to compare such signatures using Wasserstein distances.

For two probability measures P and Q over \mathbb{R}_+ , the L_1 -Wasserstein distance can be rewritten as the L_1 -norm between the cumulative distribution functions of the measures, $F_P : t \mapsto P((-\infty, t])$ and F_Q , or as well, as the L_1 -norm between the quantile functions, $F_P^{-1} : s \mapsto \inf\{x \in \mathbb{R} \mid F(x) \geq s\}$ and F_Q^{-1} . Thus, the computation of the L_1 -Wasserstein distance between empirical measures is easy, in $O(N \log(N))$ for two empirical measures from subsets of \mathbb{R} of size N , the complexity of a sort.

Shape signatures are widely used for classification or pre-classification tasks; see for instance [OFCD02]. With a more topological point of view, persistence diagrams have been used for this purpose in [CCSG⁺09, CDSO14]. But, as far as we know, the construction of well-founded statistical tests from signatures to compare mm-spaces has not been considered among the literature, this is the purpose of Section 2.2.

In the following, we will study stable and discriminative properties of the DTM-signature, that is a new signature that derives from the distance-to-measure function, as follows.

Definition 2.11 (DTM-signature). The *DTM-signature* associated to some mm-space (\mathcal{X}, δ, P) , denoted $d_{P,h}(P)$, is the distribution of the real-valued random variable $d_{P,h}(X)$ where X is some random variable of distribution P .

Study of a stable and discriminative descriptor of measures up to an isomorphism – The DTM-signature

Stability of the DTM-signature In this section, we prove stability results for the DTM-signature. These results all rely on the stability of the distance-to-measure function itself.

The DTM-signature turns out to be stable with respect to this Gromov-Wasserstein distance.

Proposition 2.12. *We have that:*

$$W_1(d_{P,h}(P), d_{Q,h}(Q)) \leq \frac{1}{h} GW(\mathcal{X}, \mathcal{Y}).$$

The proof of Proposition 2.12 is to be found in Section 2.3.1. The proof is relatively similar to the ones given by Mémoli in [Mém11] for other signatures.

It follows directly that two isomorphic mm-spaces have the same DTM-signature. Whenever the two mm-spaces are *embedded*⁹ into the same space, we also get stability with respect to the L_1 -Wasserstein distance.

⁹We say that two mm-spaces (\mathcal{X}, δ, P) and (\mathcal{Y}, γ, Q) are embedded into the same space when \mathcal{X} and \mathcal{Y} are included in the same metric space $(\mathcal{Z}, \delta_{\mathcal{Z}})$ with $\delta = \delta_{\mathcal{Z}}$ on \mathcal{X} and $\gamma = \delta_{\mathcal{Z}}$ on \mathcal{Y}

Proposition 2.13. *If (\mathcal{X}, δ, P) and (\mathcal{Y}, δ, Q) are two metric spaces embedded into some metric space (\mathcal{Z}, δ) , then we can bound $W_1(d_{P,h}(P), d_{Q,h}(Q))$ above by*

$$W_1(P, Q) + \min \left\{ \|d_{P,h} - d_{Q,h}\|_{\infty, \text{Supp}(P)}, \|d_{P,h} - d_{Q,h}\|_{\infty, \text{Supp}(Q)} \right\},$$

and more generally by

$$\left(1 + \frac{1}{h}\right) W_1(P, Q).$$

The proof of Proposition 2.13 is to be found in Section 2.3.1.

As a consequence, just as the distance-to-measure, the distance-to-measure signature turns out to be stable under sampling. Moreover, set $Q = \alpha P * \mathcal{N} + (1 - \alpha)\mathcal{C}$ for \mathcal{N} and \mathcal{C} distributions. Then, provided that \mathcal{N} has a low first moment $\mathcal{N}\|u\|$, α is close to 1 and \mathcal{C} has a bounded support, the DTM-signature $d_{Q,h}(Q)$ will be close to $d_{P,h}(P)$.

Discriminative properties of the DTM-signatures The DTM-signature is stable but unfortunately does not always discriminate between mm-spaces. Indeed, in the following counterexample from [Mém11] (example 5.6), there are two non-isomorphic mm-spaces sharing the same signatures for all values of h .

Example 2.14. We consider two graphs made of 9 vertices each, clustered in three groups of 3 vertices, such that each vertex is at distance 1 exactly from each vertex of its group and at distance 2 from any other vertex. We assign a mass to each vertex; the distribution is the following, for the first graph:

$$P = \left\{ \left(\frac{23}{140}, \frac{1}{105}, \frac{67}{420} \right), \left(\frac{3}{28}, \frac{1}{28}, \frac{4}{21} \right), \left(\frac{2}{15}, \frac{1}{15}, \frac{2}{15} \right) \right\},$$

and for the second graph:

$$Q = \left\{ \left(\frac{3}{28}, \frac{1}{15}, \frac{67}{420} \right), \left(\frac{2}{15}, \frac{4}{21}, \frac{1}{105} \right), \left(\frac{23}{140}, \frac{2}{15}, \frac{1}{28} \right) \right\}.$$

The mm-spaces ensuing are not isomorphic since any one-to-one and onto measure-preserving map would send at least one pair of vertices at distance 1 from each other to a pair of vertices at distance 2 from each other, and thus it would not be an isometry.

Moreover, note that the DTM-signatures associated to the graphs are equal since the total mass of each cluster is exactly equal to $\frac{1}{3}$.

Nevertheless, the DTM-signature can be discriminative in some cases. As an example, in the following we derive lower bounds for the L_1 -Wasserstein distance between two DTM-signatures under three different alternatives.

When the distances are multiplied by some positive real number λ . Let λ be some positive real number. The DTM-signature discriminates between two mm-spaces isomorphic up to a dilatation of parameter λ , for $\lambda \neq 1$.

Proposition 2.15. *Let (\mathcal{X}, δ, P) and $(\mathcal{Y}, \gamma, Q) = (\mathcal{X}, \lambda\delta, P)$ be two mm-spaces. We have*

$$W_1(d_{P,h}(P), d_{Q,h}(Q)) = |1 - \lambda| P d_{P,h}(u).$$




Figure 2.6: P




Figure 2.7: Q

The proof of Proposition 2.15 is straightforward, see Section 2.3.1.

The case of uniform measures on non-empty bounded open subsets of \mathbb{R}^d

The DTM-signature discriminates between two uniform measures over two non-empty bounded open subsets of \mathbb{R}^d with different Lebesgue volume, provided that h is small enough. We use the same notation as in Section 2.1.1.

Proposition 2.16. *Let $(O, \|\cdot\|_2, P_O)$ and $(O', \|\cdot\|_2, P_{O'})$ be two mm-spaces, for O and O' two non-empty bounded open subsets of \mathbb{R}^d satisfying $O = \overline{(O)}^\circ$ and $O' = \overline{(O')}^\circ$.*

A lower bound for $W_1(d_{P_O, h}(P_O), d_{P_{O'}, h}(P_{O'}))$ is given by

$$\min(P_O(O_{\epsilon(h, O)}), P_{O'}(O'_{\epsilon(h, O')})) \frac{d}{d+1} \left(\frac{h}{\omega_d}\right)^{\frac{1}{d}} \left| \text{Leb}_d(O)^{\frac{1}{d}} - \text{Leb}_d(O')^{\frac{1}{d}} \right|.$$

The proof of Proposition 2.16 is deferred to Section 2.3.1. This proposition can be applied to a simple case.

Example 2.17. Let O be the open unit ball in \mathbb{R}^d , and O' the hypercube of radius 1 in \mathbb{R}^d . Then, the DTM-signature discriminates between P_O and $P_{O'}$ whenever

$$h < \frac{\omega_d}{2^d}.$$

The proof of the statement in Example 2.17 is to be found in Section 2.3.1.

The case of two measures on the same open subset of \mathbb{R}^d with one measure uniform

Let $(O, \|\cdot\|_2, P_O)$ and $(O, \|\cdot\|_2, Q)$ be two mm-spaces with O a non-empty bounded open subset of \mathbb{R}^d , $P_O = \frac{\text{Leb}_d(\cdot \cap O)}{\text{Leb}_d(O)}$ and Q a measure absolutely continuous with respect to P_O . According to the Radon-Nikodym theorem, there is some P_O -measurable function f on O such that, for all Borel sets A in O ,

$$Q(A) = \int_A f(\omega) dP_O(\omega).$$

We can consider the λ -super-level sets of the function f denoted by $\{f \geq \lambda\}$. Again, we will denote by $\{f \geq \lambda\}_\epsilon$ the set of points belonging to $\{f \geq \lambda\}$ whose distance to $\partial\{f \geq \lambda\}$ is at least ϵ .

Then we get the following lower bound for the L_1 -Wasserstein distance between the two signatures:

Proposition 2.18. *Under these hypotheses, a lower bound for $W_1(d_{P_O,h}(P_O), d_{Q,h}(Q))$ is given by*

$$\frac{1}{1+d} \frac{1}{\text{Leb}_d(O)} \left(\frac{h \text{Leb}_d(O)}{\omega_d} \right)^{\frac{1}{d}} \int_{\lambda=1}^{\infty} \frac{1}{\lambda^{\frac{1}{d}}} \max_{\lambda' \geq \lambda} \text{Leb}_d \left(\{f \geq \lambda'\} \left(\frac{h \text{Leb}_d(O)}{\omega_d \lambda'} \right)^{\frac{1}{d}} \right) d\lambda,$$

where ω_d stands for $\text{Leb}_d(\mathcal{B}(0, 1))$, the Lebesgue volume of the unit d -dimensional ball.

The proof of Proposition 2.18 is to be found in Section 2.3.1.

When the density f is Hölder

In the following, we assume that $\text{Reach}(O) > 0$ and that f is Hölder on O , with positive parameters $\chi \in (0, 1]$ and $L > 0$, that is:

$$\forall x, y \in O, |f(x) - f(y)| \leq L \|x - y\|_2^\chi.$$

Then for h small enough, the DTM-signature is discriminative.

Proposition 2.19. *Under the previous assumptions, if one of the following conditions is satisfied, then the quantity $W_1(d_{P_O,h}(P_O), d_{Q,h}(Q))$ is positive:*

$$h < \frac{\omega_d}{\text{Leb}_d(O)} \min \left\{ \text{Reach}(O)^d, \left(\frac{\|f\|_{\infty, O} - 1}{2L} \right)^{\frac{d}{\chi}} \right\};$$

$$h \in \left[\frac{\omega_d}{\text{Leb}_d(O)} (\text{Reach}(O))^d, (\|f\|_{\infty, O} - 2L(\text{Reach}(O))^\chi) (\text{Reach}(O))^d \frac{\omega_d}{\text{Leb}_d(O)} \right);$$

$$h \in \left[\frac{\omega_d}{\text{Leb}_d(O)} \left(\frac{d}{\chi} \right)^{\frac{d}{\chi}} (2L)^{-\frac{d}{\chi}}, \min \left\{ h_0, \frac{\omega_d}{\text{Leb}_d(O)} (\text{Reach}(O))^{d+\chi} \frac{\chi}{d} 2L \right\} \right),$$

with $h_0 = \|f\|_{\infty, O}^{\frac{d}{\chi}+1} \frac{\omega_d}{\text{Leb}_d(O)} \left(\frac{d}{\chi} \right)^{\frac{d}{\chi}} (2L)^{-\frac{d}{\chi}} \left(\frac{\chi}{d+\chi} \right)^{\frac{\chi}{d+\chi}}$.

Moreover, under any of these conditions, we get the lower bound for $W_1(d_{P_O,h}(P_O), d_{Q,h}(Q))$, our pseudo-metric:

$$\frac{1}{1+d} \left(\frac{h \text{Leb}_d(O)}{\omega_d} \right)^{\frac{1}{d}} \int_{\lambda=1}^{\infty} \frac{1}{\lambda^{1+\frac{1}{d}}} \sup_{\lambda' \geq \lambda} Q(\{f \geq \lambda' + L\epsilon(\lambda')^\chi\} \cap O_{\epsilon(\lambda')}) d\lambda,$$

with $\epsilon(\lambda') = \lambda'^{-\frac{1}{d}} \left(\frac{h \text{Leb}_d(O)}{\omega_d} \right)^{\frac{1}{d}}$.

Here, ω_d stands for $\text{Leb}_d(\mathcal{B}(0, 1))$, the Lebesgue volume of the unit d -dimensional ball.

The proof of Proposition 2.19 is to be found in Section 2.3.1.

This proposition displays different intervals of values for h for which the DTM-signatures are discriminative. These intervals depend on the reach of O . Indeed, for h small enough (i.e. smaller than a function of $\text{Reach}(O)$), Proposition 2.6 entails that the DTM of P_O is known for most of the elements of O . For such elements, it will be easier to derive bounds for the difference between the DTM of P_O and the DTM of Q .

The Hölder hypothesis implies the DTM of Q at two points close enough in O are close. The bounds obtained come from the simple idea that the density of Q is not constant. Then, for some h , we can exhibit a set of points with positive P_O -measure for which the DTM of Q is smaller than the minimum of the DTM of P_O .

Note that these intervals do not necessarily represent all of the values of h for which the distance-to-measure is discriminative. Furthermore, they do not always correspond to the most discriminative values of h . At least, this proposition can be applied to concrete cases, proving the existence of some mass parameters h for which the DTM-signature is discriminative.

Example 2.20. Let O be the open unit ball $\mathcal{B}(0, 1)$ in \mathbb{R}^d , P_O the uniform measure on O and Q the multivariate normal distribution $N(0, \sigma^2 I)$ restricted to $\mathcal{B}(0, 1)$.

The signatures are discriminative, provided that σ is small enough, for all h smaller than 0.23 if $d = 1$; 0.30 if $d = 2$; 0.68 if $d = 3$; and for all values of h if $d \geq 4$.

The proof of the statement in Example 2.20 is to be found in Section 2.3.1.

In Section 2.2, we will build a statistical test of isomorphism based on the DTM-signature. Morally, whenever the DTM-signatures are discriminative, the test will be powerful. The previous examples provide several relevant examples for which the DTM-signature turns out to be discriminative. Thus, for these examples among others, the test will be powerful.

2.2 Statistical test of similarity in the Gromov-Wasserstein sense

Let (\mathcal{X}, δ, P) and (\mathcal{Y}, γ, Q) be two mm-spaces.

In this section, we build statistical tests of the null hypothesis

$$H_0 \text{ "The mm-spaces } (\mathcal{X}, \delta, P) \text{ and } (\mathcal{Y}, \gamma, Q) \text{ are isomorphic" ,}$$

against its alternative:

$$H_1 \text{ "The mm-spaces } (\mathcal{X}, \delta, P) \text{ and } (\mathcal{Y}, \gamma, Q) \text{ are not isomorphic" .}$$

The statistical tests we propose are based on the observation of two samples, an N -sample from P and an N' -sample from Q . To simplify notation, we assume that $N' = N$, but the methods proposed also work when N' is different from N . The important thing is to keep the same n in both cases, as defined below.

Given an N -sample X_1, X_2, \dots, X_N from the measure P , we denote $\hat{P}_N = \frac{1}{N} \sum_{i=1}^N \delta_{X_i}$ and $\hat{P}_n = \frac{1}{n} \sum_{i=1}^n \delta_{X_i}$ for some $n \leq N$. As well, we define \hat{Q}_N and \hat{Q}_n from an N -sample from Q .

We recall that $d_{\hat{P}_N, h}(\hat{P}_n)$ is the discrete distribution $\frac{1}{n} \sum_{i=1}^n \delta_{d_{\hat{P}_N, h}(X_i)}$, and that we compare signatures with W_1 , the L_1 -Wasserstein distance.

The *test statistic* is then defined as

$$T_{N, n, h}(P, Q) = \sqrt{n} W_1 \left(d_{\hat{P}_N, h}(\hat{P}_n), d_{\hat{Q}_N, h}(\hat{Q}_n) \right),$$

and its distribution is denoted by $\mathcal{L}_{N, n, h}(P, Q)$.

Note that for two isomorphic mm-spaces (\mathcal{X}, δ, P) and (\mathcal{Y}, γ, Q) , the following distributions $\mathcal{L}_{N, n, h}(P, P)$, $\mathcal{L}_{N, n, h}(Q, Q)$, and $\frac{1}{2} \mathcal{L}_{N, n, h}(P, P) + \frac{1}{2} \mathcal{L}_{N, n, h}(Q, Q)$ are equal. They correspond to the distribution of the test statistic $T_{N, n, h}(P, Q)$.

Lemma 2.21 (EQUALITY OF EMPIRICAL SIGNATURES UNDER THE ISOMORPHIC ASSUMPTION). *If (\mathcal{X}, δ, P) and (\mathcal{Y}, γ, Q) are two isomorphic mm-spaces, then the distributions of the random variables*

$$\sqrt{n}W_1(d_{\hat{P}_{N,h}}(\hat{P}_n), d_{\hat{P}'_{N,h}}(\hat{P}'_n))$$

and

$$\sqrt{n}W_1(d_{\hat{P}_{N,h}}(\hat{P}_n), d_{\hat{Q}_{N,h}}(\hat{Q}_n))$$

are equal. Here the empirical measures are all independent and the measures \hat{P}'_N and \hat{P}'_n are from samples from P .

The proof of Lemma 2.21 is to be found in Section 2.3.2

For some $\alpha > 0$, we denote by $q_\alpha = \inf\{x \in \mathbb{R} \mid F(x) \geq 1 - \alpha\}$, the α -quantile of a distribution with cumulative distribution function F .

The α -quantile $q_{\alpha, N, n, h}$ of $\frac{1}{2}\mathcal{L}_{N, n, h}(P, P) + \frac{1}{2}\mathcal{L}_{N, n, h}(Q, Q)$ will be approximated by the α -quantile $\hat{q}_{\alpha, N, n, h}$ of $\frac{1}{2}\mathcal{L}_{N, n, h}^*(\hat{P}_N, \hat{P}_N) + \frac{1}{2}\mathcal{L}_{N, n, h}^*(\hat{Q}_N, \hat{Q}_N)$. Here $\mathcal{L}_{N, n, h}^*(\hat{P}_N, \hat{P}_N)$ stands for the distribution of $\sqrt{n}W_1(d_{\hat{P}_{N,h}}(P_n^*), d_{\hat{P}_{N,h}}(P_n'^*))$ conditionally to \hat{P}_N , where P_n^* and $P_n'^*$ are two empirical measures from independent n -samples from \hat{P}_N .

The test we deal with in this paper is then

$$\phi_{N, n, h} = \mathbb{1}_{T_{N, n, h}(P, Q) \geq \hat{q}_{\alpha, N, n, h}}.$$

The null hypothesis H_0 is rejected if $\phi_{N, n, h} = 1$, that is if the L_1 -Wasserstein distance between the two empirical signatures $d_{\hat{P}_{N,h}}(\hat{P}_n)$ and $d_{\hat{Q}_{N,h}}(\hat{Q}_n)$ is too high.

Note that it is equivalent to compute a p -value $\hat{p}_{N, n, h}$ from the subsampling distribution and the test statistic:

$$\hat{p}_{N, n, h} = 1 - F^*(T_{N, n, h}(P, Q)), \quad (2.3)$$

with $F^*(t) = \mathbb{P}(T \leq t)$ for T a random variable from the distribution $\frac{1}{2}\mathcal{L}_{N, n, h}^*(\hat{P}_N, \hat{P}_N) + \frac{1}{2}\mathcal{L}_{N, n, h}^*(\hat{Q}_N, \hat{Q}_N)$.

The statistical test consists in rejecting the hypothesis H_0 if the p -value is not larger than α , that is

$$\phi_{N, n, h} = \mathbb{1}_{\hat{p}_{N, n, h} \leq \alpha}.$$

2.2.1 A test of asymptotic level α

In this section, we derive assumptions under which the test we propose is of asymptotic level α , that is such that

$$\limsup_{N \rightarrow \infty} \mathbb{P}_{(P, Q) \sim H_0}(\phi_{N, n, h} = 1) \leq \alpha,$$

where $\mathbb{P}_{(P, Q) \sim H_0}(\phi_{N, n, h} = 1)$ stands for the probability that $\phi_{N, n, h} = 1$ when the test is build from samples from two mm-spaces (\mathcal{X}, δ, P) and (\mathcal{Y}, γ, Q) which are isomorphic.

In the following, $\mathbb{G}_{P, h}$ and $\mathbb{G}'_{P, h}$ represent two independent Gaussian processes with covariance kernel $\kappa(s, t) = F_{d_{P, h}(P)}(s) (1 - F_{d_{P, h}(P)}(t))$ for $s \leq t$, and $\|\mathbb{G}_{P, h} - \mathbb{G}'_{P, h}\|_1$ denotes the random variable corresponding to the integral on $[0, 1]$ of the absolute value of the difference

between the two Gaussian processes.

This first theorem describes the case of measures supported on compact subsets of \mathbb{R}^d equipped with the Euclidean metric. Better results are obtained for measures which are more regular in the following sense. We say that a measure P is (a, b) -standard with positive parameters a and b , if for any positive radius r and any point x of the support of P , we have that $P(B(x, r)) \geq \min\{1, ar^b\}$, with $B(x, r) = \{y \in \mathcal{X} \mid \delta(x, y) < r\}$.

Theorem 2.22. *Let P and Q be two Borel probability measures supported on compact subsets of \mathbb{R}^d . We set $N = cn^\rho$ with some $c > 0$ and $\rho > 1$.*

The statistical test

$$\phi_{N,n,h} = \mathbb{1}_{\sqrt{n}W_1(d_{\hat{P}_{N,h}}(\hat{P}_n), d_{\hat{Q}_{N,h}}(\hat{Q}_n)) \geq \hat{q}_{\alpha, N, n, h}}$$

is of asymptotic level α

- *in the general case, if $\rho > \frac{\max\{d, 2\}}{2}$,*
- *in the (a, b) -standard case, if $\rho > 1$,*

with the additional assumption that $\mathcal{L}(\|\mathbb{G}_{P,h} - \mathbb{G}'_{P,h}\|_1)$ is atomless.

This second theorem describes the case of measures supported on compact subsets of general metric spaces, but requires assumptions that are more difficult to check.

Theorem 2.23. *Let P and Q be two Borel probability measures supported on compact subsets of general metric spaces.*

Let n be chosen as a function of N such that when N goes to infinity, n goes to infinity and $\sqrt{n}\mathbb{E}[\|\mathbb{d}_{P,h} - \mathbb{d}_{\hat{P}_{N,h}}\|_{\infty, \mathcal{X}}]$ goes to zero. Assume moreover that $\sqrt{n}W_1(\mathbb{d}_{P,h}(P), \mathbb{d}_{P,h}(\hat{P}_N))$ and $\sqrt{n}\|\mathbb{d}_{P,h} - \mathbb{d}_{\hat{P}_{N,h}}\|_{\infty, \mathcal{X}}$ go to zero a.e..

Then, if $\mathcal{L}(\|\mathbb{G}_{P,h} - \mathbb{G}'_{P,h}\|_1)$ is atomless, the statistical test

$$\phi_{N,n,h} = \mathbb{1}_{\sqrt{n}W_1(\mathbb{d}_{\hat{P}_{N,h}}(\hat{P}_n), \mathbb{d}_{\hat{Q}_{N,h}}(\hat{Q}_n)) \geq \hat{q}_{\alpha, N, n, h}}$$

is of asymptotic level α .

Note that for uniform measures on any sphere in \mathbb{R}^d , the continuity assumption for the cumulative distribution function of $\mathcal{L}(\|\mathbb{G}_{P,h} - \mathbb{G}'_{P,h}\|_1)$ is not satisfied. This is a degenerate case. Thus, the test cannot be applied to such mm-spaces directly. Nonetheless, it can be applied after adding isotropic noise to both samples.

These theorems are obtained by proving that the distribution of the test statistic

$$\frac{1}{2}\mathcal{L}_{N,n,h}(P, P) + \frac{1}{2}\mathcal{L}_{N,n,h}(Q, Q)$$

under the hypothesis H_0 and the subsampling distribution

$$\frac{1}{2}\mathcal{L}_{N,n,h}^*(\hat{P}_N, \hat{P}_N) + \frac{1}{2}\mathcal{L}_{N,n,h}^*(\hat{Q}_N, \hat{Q}_N)$$

converge weakly to the fixed distribution

$$\frac{1}{2}\mathcal{L}(\|\mathbb{G}_{P,h} - \mathbb{G}'_{P,h}\|_1) + \frac{1}{2}\mathcal{L}(\|\mathbb{G}_{Q,h} - \mathbb{G}'_{Q,h}\|_1)$$

when n and N go to ∞ .

In order to adopt a non-asymptotic and more visual point of view, we also derive upper bounds in expectation for the L_1 -Wasserstein distance between the first two distributions above.

Note that it is sufficient to prove weak convergence for $\mathcal{L}_{N,n,h}(P, P)$ and $\mathcal{L}_{N,n,h}^*(\hat{P}_N, \hat{P}_N)$. Moreover,

$$W_1 \left(\frac{1}{2} \mathcal{L}_{N,n,h}(P, P) + \frac{1}{2} \mathcal{L}_{N,n,h}(Q, Q), \frac{1}{2} \mathcal{L}_{N,n,h}^*(\hat{P}_N, \hat{P}_N) + \frac{1}{2} \mathcal{L}_{N,n,h}^*(\hat{Q}_N, \hat{Q}_N) \right)$$

is bounded above by

$$\frac{1}{2} W_1 \left(\mathcal{L}_{N,n,h}(P, P), \mathcal{L}_{N,n,h}^*(\hat{P}_N, \hat{P}_N) \right) + \frac{1}{2} W_1 \left(\mathcal{L}_{N,n,h}(Q, Q), \mathcal{L}_{N,n,h}^*(\hat{Q}_N, \hat{Q}_N) \right).$$

This is a straightforward consequence of the definition of the L_1 -Wasserstein distance with transport plans. Thus, this is also sufficient to derive upper bounds in expectation for the quantity $W_1 \left(\mathcal{L}_{N,n,h}(P, P), \mathcal{L}_{N,n,h}^*(\hat{P}_N, \hat{P}_N) \right)$.

Theorem 2.23 follows from the following two lemmas.

Lemma 2.24. *For P a measure supported on a compact set, we choose n as a function of N such that: when N goes to infinity, n goes to infinity, $\sqrt{n} \mathbb{E}[\|d_{P,h} - d_{\hat{P}_N,h}\|_{\infty, \mathcal{X}}]$ goes to zero or more specifically $\frac{\sqrt{n}}{h} \mathbb{E}[W_1(P, \hat{P}_N)]$ goes to zero. Then we have that*

$$\mathcal{L}_{N,n,h}(P, P) \rightsquigarrow \mathcal{L}(\|\mathbb{G}_{P,h} - \mathbb{G}'_{P,h}\|_1),$$

when N goes to infinity. Moreover, if n is chosen such that $\sqrt{n} W_1(d_{P,h}(P), d_{P,h}(\hat{P}_N))$ and $\sqrt{n} \|d_{P,h} - d_{\hat{P}_N,h}\|_{\infty, \mathcal{X}}$ go to zero a.e., we have that for almost every sample $X_1, X_2, \dots, X_N, \dots$,

$$\mathcal{L}_{N,n,h}^*(\hat{P}_N, \hat{P}_N) \rightsquigarrow \mathcal{L}(\|\mathbb{G}_{P,h} - \mathbb{G}'_{P,h}\|_1),$$

when N goes to infinity.

The proof of Lemma 2.24 is to be found in Section 2.3.2.

Lemma 2.25. *If the two weak convergences in lemma 2.24 occur, and if the α -quantile q_α of the distribution $\mathcal{L}(\frac{1}{2} \|\mathbb{G}_{P,h} - \mathbb{G}'_{P,h}\|_1 + \frac{1}{2} \|\mathbb{G}_{Q,h} - \mathbb{G}'_{Q,h}\|_1)$ is a point of continuity of its cumulative distribution function, then the asymptotic level of the test at (P, Q) is α .*

The proof of Lemma 2.25 is to be found in Section 2.3.2.

Theorem 2.22 follows from Theorem 2.23 and some rates of convergence for the Wasserstein metric between a measure and its empirical version in [FG15], or for the infinity norm of the difference between the distance-to-measure and the distance to the empirical measure in [CMM16]. These results are recalled and proved in Proposition 2.26 and Proposition 2.28 with complementary results.

The case of measures supported on a compact subset of \mathbb{R}^d

We choose $N = cn^\rho$ for some positive constants ρ and c . Then the test is asymptotically valid for two measures supported on a compact subset of the Euclidean space \mathbb{R}^d if we assume that $\rho > \frac{\max\{d, 2\}}{2}$.

Proposition 2.26. *Let P be some Borel probability measure supported on some compact subset of \mathbb{R}^d . Under the assumption*

$$\rho > \frac{\max\{d, 2\}}{2},$$

the two weak convergences of lemma 2.24 occur.

Moreover, a bound for the expectation of $W_1\left(\mathcal{L}_{N,n,h}(P, P), \mathcal{L}_{N,n,h}^(\hat{P}_N, \hat{P}_N)\right)$ is of order:*

$$N^{\frac{1}{2\rho} - \frac{1}{\max\{d, 2\}}} (\log(1 + N))^{\mathbb{1}_{d=2}}.$$

Furthermore, $W_1\left(\mathcal{L}_{N,n,h}(P, P), \mathcal{L}_{N,n,h}^(\hat{P}_N, \hat{P}_N)\right) \rightarrow 0$ a.e. when n goes to ∞ .*

The proof of Proposition 2.26 is deferred to Section 2.3.2.

The case of (a, b) -standard measures supported on a compact subset of \mathbb{R}^d

Recall that a probability measure P is (a, b) -standard with positive parameters a and b , if for any positive radius r and any point x of the support of P , we have that $P(\mathcal{B}(x, r)) \geq \min\{1, ar^b\}$. Uniform measures on open subsets of \mathbb{R}^d satisfy such a property:

Example 2.27. Let O be a non-empty bounded open subset of \mathbb{R}^d . Then, the measure P_O is (a, d) -standard with

$$a = \frac{\omega_d}{\text{Leb}_d(O)} \left(\frac{\text{Reach}(O)}{\mathcal{D}(O)} \right)^d.$$

Here, $\mathcal{D}(O)$ stands for the diameter of O , ω_d for $\text{Leb}_d(\mathcal{B}(0, 1))$, the Lebesgue volume of the unit d -dimensional ball, and $\text{Reach}(O)$ is the reach of the open set O as defined in Section 2.1.2.

The proof of the statement in Example 2.27 is to be found in Section 2.3.2.

Similar results can be obtained for uniform measures on compact submanifolds of dimension d . In [NSW08] (lemma 5.3), the authors give a bound for a depending on the reach of the sub-manifold.

The test is asymptotically valid for two (a, b) -standard measures supported on compact connected subsets of \mathbb{R}^d if $\rho > 1$:

Proposition 2.28. *Let P be an (a, b) -standard measure supported on a connected compact subset of \mathbb{R}^d . The two weak convergences of lemma 2.24 occur if the assumption $\rho > 1$ is satisfied. Moreover, a bound for the expectation of $W_1\left(\mathcal{L}_{N,n,h}(P, P), \mathcal{L}_{N,n,h}^*(\hat{P}_N, \hat{P}_N)\right)$ is of order $N^{\frac{1}{2\rho} - \frac{1}{2}}$ up to a logarithmic term.*

The proof of Proposition 2.28 is to be found in Section 2.3.2.

Note that we can achieve a rate close to the parametric rate for standard measures, whereas for general measures, the rate gets worse when the dimension increases. In any case, we need ρ to be as large as possible for the subsampling distribution to be a good enough approximation of the distribution of the statistic, that is to have a type I error close enough to α , keeping in mind that n should go to ∞ with N .

2.2.2 The power of the test

The power of the test $\phi_{N,n,h} = \mathbb{1}_{\sqrt{n}W_1(d_{\hat{P}_N,h}(\hat{P}_n), d_{\hat{Q}_N,h}(\hat{Q}_n)) \geq \hat{q}_{\alpha,N,n,h}}$ is defined for two mm-spaces (\mathcal{X}, δ, P) and (\mathcal{Y}, γ, Q) by

$$1 - \mathbb{P}_{(P,Q)}(\phi_{N,n,h} = 0),$$

where $\mathbb{P}_{(P,Q)}(\phi_{N,n,h} = 0)$ stands for the probability that $\phi_{N,n,h} = 0$ when the test is build from samples from two general mm-spaces (\mathcal{X}, δ, P) and (\mathcal{Y}, γ, Q) .

If the spaces are not isomorphic, we want the test to reject H_0 with high probability. It means that we want the power to be as large as possible. Here, we give a lower bound for the power, or more precisely an upper bound for $\mathbb{P}_{(P,Q)}(\phi_{N,n,h} = 0)$, the *type II error*.

Theorem 2.29. *Let P and Q be two Borel measures supported on \mathcal{X} and \mathcal{Y} , two compact subsets of \mathbb{R}^d . We assume that the mm-spaces (\mathcal{X}, δ, P) and (\mathcal{Y}, γ, Q) are non-isomorphic and that the DTM-signature is discriminative for some h in $(0, 1]$, meaning that the pseudo-metric $W_1(d_{P,h}(P), d_{Q,h}(Q))$ is positive. We choose $N = n^\rho$ with $\rho > 1$. Then for all positive ϵ , there exists n_0 depending on P and Q such that for all $n \geq n_0$, the type II error*

$$\mathbb{P}_{(P,Q)}\left(\sqrt{n}W_1\left(d_{\hat{P}_{N,h}}(\hat{P}_n), d_{\hat{Q}_{N,h}}(\hat{Q}_n)\right) < \hat{q}_{\alpha,N,n,h}\right)$$

is bounded above by

$$4 \exp\left(-\frac{W_1^2(d_{P,h}(P), d_{Q,h}(Q))}{(2 + \epsilon) \max\{\mathcal{D}_{P,h}^2, \mathcal{D}_{Q,h}^2\}} n\right),$$

with $\mathcal{D}_{P,h}$ the diameter of the support of the measure $d_{P,h}(P)$.

The proof of Theorem 2.29 is to be found in Section 2.3.2.

In order to have a high power, that is to reject H_0 more often when the mm-spaces are not isomorphic, we need n to be big enough, that is ρ small enough. Recall that n has to be small enough for the law of the statistic and its subsampling version to be close. This means that some compromise must be made. Moreover, the choice of h for the test should depend on the geometry of the mm-spaces. The tuning of these parameters from the data is still an open question.

Moreover, note that the power of the test is strongly related to $W_1(d_{P,h}(P), d_{Q,h}(Q))$. The test is powerful when the pseudo-metric is high, and does not discriminate between measures when it is low. In the following section, we derive some upper-bounds and lower-bounds for this pseudo-metric, under some geometric assumptions.

In this section, we first describe the procedure to implement the statistical test of isomorphism. Then, we illustrate the validity of the method by providing some numerical approximations of the type-I error and the power of the test for various examples. We also compare our test to a more basic statistical test of isomorphism.

2.2.3 The algorithm

The procedure for the statistical test is as follows.

In the code, if $\mathbb{Z} = \{Z_1, Z_2, \dots, Z_n\}$, then we use the notation $\mathbb{1}_{\mathbb{Z}}$ for the measure $\frac{1}{n} \sum_{i=1}^n \delta_{Z_i}$.

Algorithm 2.1: Test Procedure

Input :

$\mathbb{X} = \{X_1, X_2, \dots, X_N\}$ N -sample from P ;

$\mathbb{Y} = \{Y_1, Y_2, \dots, Y_{N'}\}$ N' -sample from Q ;

parameter n , mass parameter h , level α , number of subsampling repetitions N_{sub} ;

Compute T the test statistic

Let σ be a random permutation of $\{1, 2, \dots, N\}$;

Let σ' be a random permutation of $\{1, 2, \dots, N'\}$ independent of σ ;

```

Define  $\mathbb{X}_n = \{X_{\sigma(1)}, X_{\sigma(2)}, \dots, X_{\sigma(n)}\}$ ;
Define  $\mathbb{Y}_n = \{Y_{\sigma'(1)}, Y_{\sigma'(2)}, \dots, Y_{\sigma'(n)}\}$ ;

Define the test statistic  $T = \sqrt{n}W_1(d_{\mathbb{1}_{\mathbb{X}}, h}(\mathbb{1}_{\mathbb{X}_n}), d_{\mathbb{1}_{\mathbb{Y}}, h}(\mathbb{1}_{\mathbb{Y}_n}))$ ;

# Compute  $W_{sub}$ , a  $N_{sub}$ -sample from the subsampling law
Compute  $dtm\mathbb{X} = \{d_{\mathbb{1}_{\mathbb{X}}, h}(X_1), d_{\mathbb{1}_{\mathbb{X}}, h}(X_2), \dots, d_{\mathbb{1}_{\mathbb{X}}, h}(X_N)\}$ ;
Compute  $dtm\mathbb{Y} = \{d_{\mathbb{1}_{\mathbb{Y}}, h}(Y_1), d_{\mathbb{1}_{\mathbb{Y}}, h}(Y_2), \dots, d_{\mathbb{1}_{\mathbb{Y}}, h}(Y_{N'})\}$ ;
Let  $W_{sub}$  be empty;
for j in  $1.. \lfloor N_{sub}/2 \rfloor$ :
  Let  $dtm\mathbb{X}_1$  and  $dtm\mathbb{X}_2$  be two independent  $n$ -samples from  $dtm\mathbb{X}$  with replacement;
  Let  $dtm\mathbb{Y}_1$  and  $dtm\mathbb{Y}_2$  be two independent  $n$ -samples from  $dtm\mathbb{Y}$  with replacement;
  Add  $\sqrt{n}W_1(\mathbb{1}_{dtm\mathbb{X}_1}, \mathbb{1}_{dtm\mathbb{X}_2})$  and  $\sqrt{n}W_1(\mathbb{1}_{dtm\mathbb{Y}_1}, \mathbb{1}_{dtm\mathbb{Y}_2})$  to  $W_{sub}$ ;

# Compute  $pval$ , the p-value of the statistical test
Let  $pval$  be equal to the mean number of elements in  $W_{sub}$  bigger than  $T$ ;

Output:
The hypothesis retained is  $H_1$  if  $pval \leq 0.05$ ,  $H_0$  if not.

```

Recall that the L_1 -Wasserstein distance W_1 is simply equal to the L_1 -norm of the difference between the cumulative distribution functions, which is easy to implement in the discrete case. As explained in the Introduction, in order to compute the distance to an empirical measure on an N -sample at a point x , it is sufficient to search for its $k = \lceil hn \rceil$ -nearest neighbours in the sample, where $h \in [0, 1]$ is the mass parameter. The distance to the empirical measure can also be implemented by the R function `dtm` with tuning parameter $r = 1$, from the package TDA [FKLM14].

An example in \mathbb{R}^2

In this subsection, we will compare the statistical test of this paper (DTM) with the statistical test (KS) which consists in applying a Kolmogorov-Smirnov two-sample test to the $\frac{N}{2}$ -sample

$$\{\delta(X_1, X_2), \delta(X_3, X_4), \dots, \delta(X_{N-1}, X_N)\}$$

and the $\frac{N'}{2}$ -sample

$$\{\gamma(Y_1, Y_2), \gamma(Y_3, Y_4), \dots, \gamma(Y_{N'-1}, Y_{N'})\}$$

given an N -sample $\mathbb{X} = \{X_1, X_2, \dots, X_N\}$ from an mm-space (\mathcal{X}, δ, P) and an N' -sample $\mathbb{Y} = \{Y_1, Y_2, \dots, Y_{N'}\}$ from an mm-space (\mathcal{Y}, γ, Q) .

We apply our isomorphism test to measures supported on spirals in \mathbb{R}^2 . For some shape parameter $v \in \mathbb{R}_+$, the measure P_v is the distribution of the random vector $(R \sin(vR) + 0.03S, R \cos(vR) + 0.03S')$, with R , S and S' independent random variables, S and S' from the standard normal distribution $N(0, 1)$ and R uniform on $(0, 1)$. In the following experiments, we choose $P = P_{10}$ and $Q = P_v$ for $v \in \{15, 20, 30, 40, 100\}$.

First, from the measure $P = P_{10}$ we get an $N = 2000$ -sample $\mathbb{X} = \{X_1, X_2, \dots, X_N\}$. As well, we get an $N = 2000$ -sample $\mathbb{Y} = \{Y_1, Y_2, \dots, Y_N\}$ from the measure $Q = P_{20}$. This leads to the empirical measures \hat{P}_N and \hat{Q}_N . In Figure 2.8, we plot the cumulative distribution function of the measure $d_{\hat{P}_N, h}(\hat{P}_N)$, that is, the function F defined for all t in \mathbb{R} by the proportion of the

X_i in \mathcal{X} satisfying $d_{\hat{P}_N, h}(X_i) \leq t$. It approximates the true cumulative distribution function associated to the DTM-signature $d_{P, h}(P)$. As well, we plot the cumulative distribution function of the measure $d_{\hat{Q}_N, h}(\hat{Q}_N)$. Observe that the signatures are different. Thus, for the choice of parameter $h = 0.05$, the DTM-signature discriminates well between the measures $P = P_{10}$ and $Q = P_{20}$.

In Figure 2.9, for $h = 0.05$ and $n = 20$, we first generate 1000 independent realisations of the random variable $\sqrt{n}W_1(d_{\hat{P}_N, h}(\hat{P}_n), d_{\hat{P}'_N, h}(\hat{P}'_n))$, where \hat{P}_N and \hat{P}'_N are independent empirical measures from P_{10} , $N = 2000$, and \hat{P}_n and \hat{P}'_n are the empirical measures associated to the n first points of the samples. We plot the empirical cumulative distribution function associated to this 1000-sample.

As well, from two fixed N -samples from P_{10} , leading to two empirical distributions \hat{P}_N and \hat{P}'_N , we generate a set of $N_{sub} = 1000$ random variables $\sqrt{n}W_1(d_{\hat{P}_N, h}(P_n^*), d_{\hat{P}'_N, h}(P'^*_n))$, as explained in the Algorithm in Section 2.2.3, and we plot its cumulative distribution function. Note that the two cumulative distribution functions are close. This means that the α -quantile of the distribution of the test statistic is well approximated by the α -quantile of the subsampling distribution.





Figure 2.8: DTM-signature estimates, $h = 0.05$ Figure 2.9: Subsampling validity, $v = 10$, $h = 0.05$

In order to approximate the type-I error and the power, we repeat the procedure of test **DTM** described in Section 2.2.3 1000 times independently. At each step, we sample $N = 2000$ points from the measures $P = P_{10}$ and $Q = P_v$ to approximate the power, or twice P_v to approximate for the type-I error. We select the parameters $\alpha = 0.05$, $h = 0.05$, $n = 20$, and repeat subsampling $N_{sub} = 1000$ times. Then, we retain either H_0 or H_1 . The type-I error or power approximation is simply equal to the mean number of times the hypothesis H_0 was rejected among the 1000 independent experiments. We also approximate the power for the method **KS** after repeating this test procedure 1000 times independently. Note that by construction, the test **KS** is truly of level $\alpha = 0.05$. Figure 2.12 contains the numerical values we obtained using the R software.

v	15	20	30	40	100
type-I error DTM	0.043	0.049	0.050	0.051	0.050
power DTM	0.525	0.884	0.987	0.977	0.985
power KS	0.768	0.402	0.465	0.414	0.422




Figure 2.10: Type-I error and power approximations

It turns out that our isomorphism test **DTM** is of level close to $\alpha = 0.05$ and is powerful. For parameters $v \geq 20$, our test is even more discriminative than the test **KS**.

An example in $\mathbb{R}^{28 \times 28}$

In this subsection, we use our statistical test of isomorphism to compare the distribution of the digits "2" and the distribution of the digits "5" from the MNIST handwritten digits database. Each digit is represented by a picture of 28×28 pixels with grey levels, meaning that each digit $X = (x_1, x_2, \dots, x_{28 \times 28})$ can be seen as an element of $\mathbb{R}^{28 \times 28}$, where x_i is equal to the grey level of the i -th pixel. We equip $\mathbb{R}^{28 \times 28}$ with the Euclidean metric.




Figure 2.11: MNIST database of handwritten digits

The interest of the statistical test here is not to test whether a "2" and a "5" are isometric, but whether the distribution of the "2"s and the distribution of the "5"s (which are distributions on compact subsets of $\mathbb{R}^{28 \times 28}$) are equal up to an isomorphism. These distributions are clearly not equal, since their supports are different, but it is not clear that there is no rigid transformation between the set of "2" and the set of "5" preserving the measures, as for instance a simple permutation of the pixels.

The statistical test is based on the observation of $N = 5958$ "2" and $N' = 5421$ "5". In order to prove the validity of the test, we repeat 1000 times the experiment consisting in randomly splitting the set of "2" in two parts and applying the statistical test to these two samples. The type-I error approximation is equal to the mean number of times the hypothesis H_0 was rejected. We do the same with the set of "5". And we repeat these experiments for different values of $n \in \{10, 20, 30, 50, 75, 100, 200\}$ and for a fixed $m = 0.1$, we repeat subsampling $N_{sub} = 1000$ times.

n	10	20	30	50	75	100	200
"2"	0.052	0.052	0.051	0.052	0.058	0.048	0.069
"5"	0.064	0.07	0.043	0.044	0.047	0.045	0.054

Figure 2.12: Type-I error approximations

These results are encouraging since they prove that the test does not discriminate between two samples of "2" (respectively, between two samples of "5") with probability 0.95. Thus, the type-I error is of order 0.05. We choose the parameter $n = 100$ to make the test between the sample of "2" and the sample of "5". We get a p-value equal to 0. It means that we reject H_0 at any level α . So we can conclude that the distribution of the "2" and the distribution of the "5" in the MNIST database are not isomorphic.

2.2.4 A discussion of other methods

The Kolmogorov-Smirnov test applied to empirical DTM-signatures

In order to test isomorphism between two mm-spaces (\mathcal{X}, δ, P) and (\mathcal{Y}, γ, Q) from two N -samples $\{X_1, X_2, \dots, X_N\}$ and $\{Y_1, Y_2, \dots, Y_N\}$, the idea of applying a Kolmogorov-Smirnov test to the set of points $\mathcal{D}_P = \{d_{\hat{P}_N, h}(X_1), d_{\hat{P}_N, h}(X_2), \dots, d_{\hat{P}_N, h}(X_N)\}$ on one hand, and the set of points $\mathcal{D}_Q = \{d_{\hat{Q}_N, h}(Y_1), d_{\hat{Q}_N, h}(Y_2), \dots, d_{\hat{Q}_N, h}(Y_N)\}$ on the other fails drastically. Indeed, using a Kolmogorov-Smirnov test requires independence in the data, which the data \mathcal{D}_P and \mathcal{D}_Q do not satisfy.

In the following figure, we have plotted the cumulative distribution functions associated to the uniform measures on \mathcal{D}_P and \mathcal{D}_Q , where P and Q are equal and defined as in Section 2.2.3, with $v = 10$. The p-value was equal to $4.10^{-6} \leq 0.05$, thus the hypothesis H_0 was rejected.




Figure 2.13: DTM-signature estimates, $h = 0.05$

We repeated the experiment 1000 times independently, and the proportion of rejected hypotheses H_0 was equal to 0.926, instead of 0.05 for a statistical test of level 0.05.

Thus, applying a Kolmogorov-Smirnov test to empirical DTM-signatures is to be avoided.

A different value of n for the test statistic and for the subsampling distribution

In [PR94], Politis and Romano propose subsampling methods consisting in approximating the distribution of a statistic with values of the statistic built from smaller subsets of the data.

For our statistical test, since the distribution of the statistic and the subsampling distribution converge weakly to the same distribution a.e. under some assumptions, see Lemma 2.24, we can imagine fixing some parameters n and l smaller than N , choosing as a test statistic

$$T = \sqrt{l}W_1 \left(d_{\hat{P}_{N,h}}(\hat{P}_l), d_{\hat{Q}_{N,h}}(\hat{Q}_l) \right)$$

and approximating its distribution with the subsampling distribution

$$\frac{1}{2}\mathcal{L}_{N,n,h}^*(\hat{P}_N, \hat{P}_N) + \frac{1}{2}\mathcal{L}_{N,n,h}^*(\hat{Q}_N, \hat{Q}_N).$$

If we do so, consider (a, b) -standard measures supported on compact subsets of \mathbb{R}^d and choose $N = n^\rho$ and $l = n^\beta$ with $1 < \beta < \rho$, then the test is asymptotically of level α , provided that the cumulative distribution function of $\|\mathbb{G}_{P,h} - \mathbb{G}'_{P,h}\|_1$ is continuous. Indeed, the hypotheses of Lemma 2.24, $\sqrt{l} \mathbb{E}[\|d_{P,h} - d_{\hat{P}_{N,h}}\|_{\infty, \mathcal{X}}] \rightarrow 0$, $\sqrt{n}W_1(d_{P,h}(P), d_{P,h}(\hat{P}_N)) \rightarrow 0$ a.e. and $\sqrt{n}\|d_{P,h} - d_{\hat{P}_{N,h}}\|_{\infty, \mathcal{X}} \rightarrow 0$ a.e., will be satisfied then. Moreover, the proof of Theorem 2.29, which provides an upper-bound for the type-II error, can be generalized to this case, leading to the upper-bound

$$4 \exp \left(- \frac{W_1^2(d_{P,h}(P), d_{Q,h}(Q))}{(2 + \epsilon) \max \{ \mathcal{D}_{P,h}^2, \mathcal{D}_{Q,h}^2 \}} n^\beta \right)$$

for the type-II error of the test $\phi_{N,n,h} = \mathbb{1}_{\sqrt{l}W_1(d_{\hat{P}_{N,h}}(\hat{P}_l), d_{\hat{Q}_{N,h}}(\hat{Q}_l)) \geq \hat{q}_{\alpha, N, n, h}}$ if n is big enough.

Note that it is a real improvement for the power.

Nonetheless, the upper-bounds for the L_1 -Wasserstein distance between the test statistic distribution and the subsampling distribution in Proposition 2.26 and Proposition 2.28 cannot be generalized easily.

The following experiments emphasize the fact that the subsampling distribution does not necessarily well approximate the distribution of the test statistic if n and l are different. For the parameters $n = 20$ and $l \in \{20, 50, 200\}$, we have repeated 1000 times the experiment consisting of computing the p-value of our statistical test from two 2000-samples on a spiral with shape parameter $v = 10$, subsampling $N_{sub} = 1000$ times and with mass parameter $h = 0.05$. We sorted these p-values and plotted the associated cumulative distribution function. In this experiment, the hypothesis H_0 is satisfied, so the p-values should be uniformly distributed. Moreover they are independent. Thus, the curve we obtained should lie close to the diagonal. This is not the case when l is too far from n .

However, we get a power equal to 1 when choosing $l = 200$ or $l = 50$ instead of $l = 20$, which is much better than 0.884 which was obtained in Section 2.2.3 from the same experiment but with $l = n = 20$.




Figure 2.14: $l = 20$




Figure 2.15: $l = 50$




Figure 2.16: $l = 200$

Such a procedure should not be used, despite the improvement of power. Indeed, we have not proved the existence of a non-asymptotic control of a distance between the distribution of the test statistic and the subsampling distribution. Moreover, the experiments emphasize that these distributions are too different to get a test of type-I error not greater than α .

The one-sample Kolmogorov-Smirnov test of uniformity applied to p-values

A major problem of the statistical test proposed in this paper is that the hypothesis retained truly depends on the arbitrary selection of the two n -samples to build the test statistic among the $\binom{N}{n}^2$ possible pairs of n -samples. Indeed, the p-value defined by (2.3) is random in the sense that different p-values can be associated to the same two N -samples. Moreover, the power is not that high because of n , which can be very small in comparison to N .

As an example, in Figure 2.17, we split an $N = 2000$ -sample $\mathbb{X} = \{X_1, X_2, \dots, X_N\}$ from the distribution P_{10} on the spiral with shape parameter $v = 10$, into $\frac{N}{n} = 100$ disjointed subsets

$$\begin{aligned} \mathbb{X}_1 &= \{X_1, X_2, \dots, X_n\}, \\ \mathbb{X}_2 &= \{X_{n+1}, X_{n+2}, \dots, X_{2n}\}, \\ &\dots \\ \mathbb{X}_{\frac{N}{n}} &= \{X_{N-n+1}, X_{N-n+2}, \dots, X_N\}, \end{aligned}$$

with $n = 20$.

As well, we split \mathbb{Y} , an $N = 2000$ -sample from P_{10} into $\frac{N}{n} = 100$ disjointed subsets $\mathbb{Y}_1, \mathbb{Y}_2, \dots, \mathbb{Y}_{\frac{N}{n}}$.

Then, with the notation in the algorithm in Section 2.2.3, we consider the p-values $\hat{p}_1, \hat{p}_2, \dots, \hat{p}_{\frac{N}{n}}$ with \hat{p}_i associated to $T_i = \sqrt{n}W_1(d_{\mathbb{X},h}(\mathbb{1}_{\mathbb{X}_i}), d_{\mathbb{Y},h}(\mathbb{1}_{\mathbb{Y}_i}))$.

Note that the $\frac{N}{n}$ p-values would be independent if we replaced $d_{\mathbb{X},h}$ by $d_{P_{10},h}$ in the computation of the test statistic, and if the subsampling distribution was replaced with the true distribution of the statistic. In practice, when N is big enough, we are close to these assumptions. Then the p-values $\hat{p}_1, \hat{p}_2, \dots, \hat{p}_{\frac{N}{n}}$ should behave like independent random variables uniformly distributed on $[0, 1]$.

In figure 2.17, we have sorted these $\frac{N}{n} = 100$ p-values, which were built after repeating subsampling $N_{sub} = 1000$ times and with mass parameter $h = 0.05$. They seem to be uniform on $[0, 1]$; indeed their associated cumulative distribution function lies close to the diagonal.

We use this randomness to propose the following method (**DTM-KS**) to improve the power of our statistical test. We apply a one-sample Kolmogorov-Smirnov test of uniformity on $[0, 1]$ to the sample $\hat{p}_1, \hat{p}_2, \dots, \hat{p}_{\frac{N}{n}}$. Then we get a p-value p_{KS} . Thanks to the previous heuristic, this p-value should be close to uniform on $[0, 1]$ if the hypothesis H_0 of the isomorphism test was satisfied, and should be small if most of the p-values $\hat{p}_1, \hat{p}_2, \dots, \hat{p}_{\frac{N}{n}}$ were small. We can hope for a power improvement.




Figure 2.17: Sorted p-values $\hat{p}_1, \hat{p}_2, \dots, \hat{p}_{\frac{N}{n}}$




Figure 2.18: Sorted p-values p_{KS} **DTM-KS**.




Figure 2.19: Sorted p-values p_{KS} **DTM-KS2**.

In Figure 2.20, we evaluate the type-I error and the power for this new method, with the same procedure as in Section 2.2.3 and the same parameters.

v	15	20	30	40	100
type-I error DTM-KS	0.186	0.131	0.096	0.076	0.074
power DTM-KS	1	1	1	1	1

Figure 2.20: Type-I error and power approximations **DTM-KS**

In Figure 2.21, we evaluate the type-I error and the power for the testing method (**DTM-KS2**) consisting in applying a one-dimensional Kolmogorov-Smirnov test of uniformity on $[0, 1]$ to the p-values $\hat{p}'_1, \hat{p}'_2, \dots, \hat{p}'_{100}$, where \hat{p}'_i is obtained from the test statistic

$$T'_i = \sqrt{n}W_1(d_{\mathbb{X},h}(\mathbb{1}_{\mathbb{X}'_i}), d_{\mathbb{Y},h}(\mathbb{1}_{\mathbb{Y}'_i}))$$

with $\mathbb{X}'_1, \mathbb{X}'_2, \dots, \mathbb{X}'_{100}$ and $\mathbb{Y}'_1, \mathbb{Y}'_2, \dots, \mathbb{Y}'_{100}$ independent n -samples without replacement from \mathbb{X} and \mathbb{Y} respectively. The procedure is the same as in Section 2.2.3 and the parameters are the same as well.

These procedures lead to major improvements for the power, but the type-I error degrades.

v	15	20	30	40	100
type-I error DTM-KS2	0.198	0.145	0.093	0.073	0.088
power DTM-KS2	1	1	1	1	1

Figure 2.21: Type-I error and power approximations **DTM-KS2**

2.2.5 On the selection of the parameters

Naive criteria

Selection of h We have introduced a family of statistical tests parametrized with a mass parameter h . According to Theorem 2.29, the power of the test $\phi_{N,n,h}$ is an increasing function of $W_1^2(d_{P,h}(P), d_{Q,h}(Q))$. This discriminative quantity is unknown, but can be approximated in practice by $W_1^2(d_{\hat{P}_N,h}(\hat{P}_N), d_{\hat{Q}_N,h}(\hat{Q}_N))$. Thus, a natural heuristic to select the mass parameter h consists in plotting the function $h \mapsto W_1^2(d_{\hat{P}_N,h}(\hat{P}_N), d_{\hat{Q}_N,h}(\hat{Q}_N))$. The parameter h for the test can be chosen by the user, as a maximizer of this function, a local maximizer or a breaking point.

The main problems of this heuristic is that it requires the intervention of the user. Indeed, selecting the maximum is not always the best choice since the dependence in h of the power is not well understood. Moreover, there is no guarantee for this procedure. Also, this method may introduce a bias in the ensuing statistical test.

Selection of n Fix h in $(0, 1)$ and α , the level for the test. The best parameter n should be as large as possible (to have a large power), with the constraint that the statistical tests have a level close to α . A naive idea to approximate the true level of the test consists in proceeding as follows. First, we split the N -sample from P into two subsets of $\frac{N}{2}$ points. Then, we apply the statistical test $\phi_{\frac{N}{2},n,h}$ to these two samples. We repeat the experiment many times and compute the proportion of tests for which H_0 was rejected. We proceed with the same method for the sample from Q . The selected parameter n is the largest n for which this proportion is close enough to α .

Again, there is no guarantee for such a method. Moreover, by using $\frac{N}{2}$ points instead of N points for the approximation of the distance-to-measure $d_{P,h}$, we loose out. We split the sample from P into two $\frac{N}{2}$ -samples. The empirical distributions are denoted by P_1 and P_2 . An n -sample from P_i is denoted by $\hat{P}_{n,i}$. Another method consists in replacing the statistic $\sqrt{n}W_1(d_{P_1,h}(\hat{P}_{n,1}), d_{P_2,h}(\hat{P}_{n,2}))$ by $\sqrt{n}W_1(d_{\hat{P}_N,h}(\hat{P}_{n,1}), d_{\hat{P}_N,h}(\hat{P}_{n,2}))$, and the subsampling distribution accordingly. Such a method may improve the approximation of the type I error of $\phi_{N,n,h}$ under H_0 . But again, this comes without guarantees. Thus, there is many work to do about the selection of n .

From Bonferroni to Benjamini-Hochberg procedures to find out discriminative parameters h

Set $(\phi_i)_{i \in I}$ a family of statistical tests with level α . The corresponding hypotheses $(H_{0,i})_{i \in I}$ satisfy $H_0 \subset H_{0,i}$ for every $i \in I$, where H_0 is called *global null hypothesis*. Given a fixed $i \in I$, “rejecting the global null H_0 whenever $H_{0,i}$ is rejected” results in a statistical test of level α . This remark does not generalize to more than one test. Indeed, given a set I of cardinality $|I| \geq 2$, the test that rejects H_0 whenever $H_{0,i}$ is rejected for some $i \in I$, does not necessarily have a type I error inferior to α .

Multiple testing procedures such as Bonferroni or Benjamini-Hochberg are designed to overcome this issue, when the p -values $(p_i)_{i \in I}$ are all independent. Firstly, a statistical test of the global null hypothesis H_0 of type I error smaller than α is derived. Secondly, when H_0 is rejected, these procedures return a subset of I , that correspond to the set of p -values that are smaller than some threshold, determined by the procedure.

For the Bonferroni procedure, the hypothesis H_0 is rejected whenever one of the p -values p_i is smaller than $\frac{\alpha}{|I|}$. For the Benjamini-Hochberg procedure, we sort the p -values $(p_i)_{i \in I}$ in ascending order $(p^{(j)})_{j \in \llbracket 1, |I| \rrbracket}$. The threshold k is defined as the largest index j such that $p^{(j)} \leq \frac{j}{|I|} \alpha$. The hypothesis H_0 is rejected whenever this set of indices is not empty.

We retain the parameters h associated to p -values smaller than $\frac{\alpha}{|I|}$ for the first method or such that $j \leq k$ for the second method. Such parameters are the most likely to correspond to tests that are discriminative for the global null.

We can apply such a procedure to the family of tests $(\phi_{N,n,h})_{h \in I}$ with I , a finite subset of $(0, 1)$. The hypothesis $H_{0,h}$ corresponds to the assumption $d_{P,h}(P) = d_{Q,h}(Q)$, meaning that the DTM-signatures are equal. Just as in Section 2.2.4, the p -values $(\hat{p}_{N,n,h})_{h \in I}$ are random and ‘‘almost independent’’, provided that they are built on disjointed sub-samples of size n .

We repeated the experiment 100 times with 2000-samples generated on the two spirals. We set $I = \{0.95/21 \times i + 0.05 \mid i \in \llbracket 1, 20 \rrbracket\}$ and $n = 50$. We have experimented both procedures and computed the mean number of rejections of the hypothesis H_0 . This number $\hat{\alpha}$ (or $1 - \hat{\beta}$) approximates the type I error α (power $1 - \beta$) when H_0 (H_1) is satisfied. When the two spirals are the same (i.e. under the hypothesis H_0), we got $\hat{\alpha} = 0.08$ for the Benjamini-Hochberg procedure and $\hat{\alpha} = 0.07$ for the Bonferroni procedure. When the two spirals are different (i.e. under the hypothesis H_1) with parameters $v = 10$ and $v = 20$, we got $1 - \hat{\beta} = 0.87$ for the Benjamini-Hochberg procedure and $1 - \hat{\beta} = 0.86$ for the Bonferroni procedure.

2.3 Proofs

2.3.1 Proofs for Section 2.1

Proof of Proposition 2.5

This is a straightforward consequence of Proposition 2.7. The proof relies on the fact that the set of points in \mathbb{R}^d minimizing the DTM of P_O is equal to $\{x \in O \mid d_{\partial O}(x) \geq \epsilon(h, O)\}$ with $\epsilon(h, O) = \left(\frac{h \text{Leb}_d(O)}{\omega_d}\right)^{\frac{1}{d}}$, providing that the set is non-empty. Then, if $\text{Reach}(O)$ is not smaller than $\epsilon(h, O)$, O equals to the set of points at distance smaller than $\epsilon(h, O)$ from $\{x \in O \mid d_{\partial O}(x) \geq \epsilon(h, O)\}$. Thus, the measure P_O can be recovered.

Proof of Proposition 2.6

The proof of Proposition 2.6 is based on the following lemma.

Lemma 2.30. *If x in \mathbb{R}^d satisfies $P_O(\mathcal{B}(x, \epsilon)) = \frac{\omega_d \epsilon^d}{\text{Leb}_d(O)}$, then $\mathcal{B}(x, \epsilon) \subset O$.*

Proof of Lemma 2.30. If x in \mathbb{R}^d satisfies $P_O(\mathcal{B}(x, \epsilon)) = \frac{\omega_d \epsilon^d}{\text{Leb}_d(O)}$, then, $\text{Leb}_d(O^c \cap \mathcal{B}(x, \epsilon)) = 0$. Assume for the sake of contradiction that the set $O^c \cap \mathcal{B}(x, \epsilon)$ is not empty. Since $(\overline{O})^\circ = O$, then the open subset $(O^c)^\circ \cap \mathcal{B}(x, \epsilon)$ of $O^c \cap \mathcal{B}(x, \epsilon)$ is not empty, thus of positive Lebesgue measure, which is absurd. So $\mathcal{B}(x, \epsilon) \subset O$. \square

Note that for all positive l smaller than h , we have:

$$\delta_{P,l}(x) \geq \left(\frac{l \text{Leb}_d(O)}{\omega_d} \right)^{\frac{1}{d}}.$$

Moreover, these inequalities are equalities for all points x in $O_{\epsilon(h,O)}$. By integrating, we get the lower bound d_{\min} for $x \mapsto d_{P,h}(x)$, and it is attained on $O_{\epsilon(h,O)}$.

Now take some point x in \mathbb{R}^d satisfying $d_{P,h}(x) = d_{\min}$. For almost all l smaller than h , we have: $\delta_{P,l}(x) = \left(\frac{l \text{Leb}_d(O)}{\omega_d} \right)^{\frac{1}{d}}$. In particular we get for these values of l that:

$$P \left(\overline{\mathcal{B}} \left(x, \left(\frac{h \text{Leb}_d(O)}{\omega_d} \right)^{\frac{1}{d}} \right) \right) > l.$$

So, $P \left(\mathcal{B} \left(x, \left(\frac{h \text{Leb}_d(O)}{\omega_d} \right)^{\frac{1}{d}} \right) \right) = h$, and according to Lemma 2.30, we get that $x \in O_{\epsilon(h,O)}$.

Proof of Proposition 2.7

The proof of Proposition 2.7 is based on the following lemma.

Lemma 2.31. *For any x in O , there exist a maximal ball for the inclusion, included in O and containing x .*

Proof of Lemma 2.31. Let us consider the class $\mathcal{S} = \{\mathcal{B}(y, r) \mid r > 0 \text{ and } x \in \mathcal{B}(y, r) \subset O\}$ of all non-empty open balls included in O and containing x . We are going to show that this class contains a maximal element by using the Zorn's lemma. For this, we need to show that the partially-ordered set \mathcal{S} is inductive, which means that any non-empty totally-ordered subclass \mathcal{T} of \mathcal{S} is bounded above by some element of \mathcal{S} . Let \mathcal{T} be a non-empty totally-ordered subclass of \mathcal{S} . Set $R = \sup\{r > 0 \mid \exists y \in O, \mathcal{B}(y, r) \in \mathcal{T}\}$ the supremum of the radii of all balls in \mathcal{T} . Since \mathcal{T} is non-empty and O is bounded, R is positive and finite. Let $(y_k)_{k \in \mathbb{N}}$ be a sequence of centres of balls in \mathcal{T} converging to a point y in \mathbb{R}^d such that the sequence of associated radii $(r_k)_{k \in \mathbb{N}}$ is non decreasing with R as a limit. Since \mathcal{T} is totally-ordered and the radii non decreasing, the union $\bigcup_{k \in \mathbb{N}} \mathcal{B}(y_k, r_k)$ is non decreasing, equal to $\mathcal{B}(y, R)$. Thus, $\mathcal{B}(y, R)$ belongs to \mathcal{S} and upper bounds \mathcal{T} . So the class \mathcal{S} is inductive and according to the Zorn's lemma, it contains a maximal element. \square

Remind that Proposition 2.6 entails that $\{x \in \mathbb{R}^d \mid d_{P,h}(x) = d_{\min}\} = O_{\epsilon(h,O)}$. Moreover, $O_{\epsilon(h,O)}^{\epsilon(h,O)} \subset O$. Assume for the sake of contradiction that the set $O \setminus O_{\epsilon(h,O)}^{\epsilon(h,O)}$ is non-empty. Take a point x in this set and consider $\mathcal{B}(x', r')$ a maximal ball containing x and included in O given by Lemma 2.31. Since $x \notin O_{\epsilon(h,O)}^{\epsilon(h,O)}$, we get that $r' < \epsilon(h, O)$. Moreover, x' belongs to $Sk(O)$ and so, according

set ∂O , $r' = d_{\partial O}(x') \geq \text{Reach}(O) \geq \epsilon(h, O)$, which is a contradiction. So, $O_{\epsilon(h,O)}^{\epsilon(h,O)} = O$.

Proof of Proposition 2.3

From the definition of $\delta_{P,h}(x)$, we have:

$$P \left(\overline{\mathcal{B}}(x, r) \right) = \inf\{h \geq 0 \mid \delta_{P,h}(x) > r\}.$$

Moreover, since $h \rightarrow \delta_{P,h}(x)$ is right-continuous, after the differentiation the distance-to-measure function with respect to h , we have:

$$h \frac{\partial}{\partial h} d_{P,h}(x) + d_{P,h}(x) = \delta_{P,h}(x).$$

Proof of Proposition 2.12

The proof is relatively similar to the ones given by Mémoli in [Mém11] for other signatures.

For any map plan π between P and Q Borel measures on (\mathcal{X}, δ) and (\mathcal{Y}, γ) , we get:

$$\begin{aligned}
W_1(d_{P,h}(P), d_{Q,h}(Q)) &\leq \\
&\int_{\mathcal{X} \times \mathcal{Y}} |d_{P,h}(x) - d_{Q,h}(y)| d\pi(x, y) = \\
&\int_{\mathcal{X} \times \mathcal{Y}} \left| \frac{1}{h} \int_0^h \delta_{P,l}(x) dl - \frac{1}{h} \int_0^h \delta_{Q,l}(y) dl \right| d\pi(x, y) \leq \\
&\int_{\mathcal{X} \times \mathcal{Y}} \frac{1}{h} \int_0^h |\delta_{P,l}(x) - \delta_{Q,l}(y)| dl d\pi(x, y) = \\
&\frac{1}{h} \int_{\mathcal{X} \times \mathcal{Y}} \int_0^h |\inf\{r > 0 \mid P(\overline{\mathcal{B}}(x, r)) > l\} - \inf\{r > 0 \mid Q(\overline{\mathcal{B}}(y, r)) > l\}| dl d\pi(x, y) = \\
&\frac{1}{h} \int_{\mathcal{X} \times \mathcal{Y}} \int_0^h \left| \int_0^{+\infty} (\mathbb{1}_{P(\overline{\mathcal{B}}(x, r)) \leq l} - \mathbb{1}_{Q(\overline{\mathcal{B}}(y, r)) \leq l}) dr \right| dl d\pi(x, y) \leq \\
&\frac{1}{h} \int_{\mathcal{X} \times \mathcal{Y}} \int_0^{+\infty} \int_0^h |\mathbb{1}_{P(\overline{\mathcal{B}}(x, r)) \leq l} - \mathbb{1}_{Q(\overline{\mathcal{B}}(y, r)) \leq l}| dl dr d\pi(x, y) \leq \\
&\frac{1}{h} \int_{\mathcal{X} \times \mathcal{Y}} \int_0^{+\infty} |P(\overline{\mathcal{B}}(x, r)) \wedge h - Q(\overline{\mathcal{B}}(y, r)) \wedge h| dr d\pi(x, y) \leq \\
&\frac{1}{h} \int_{\mathcal{X} \times \mathcal{Y}} \int_0^{+\infty} \left| \int_{\mathcal{X} \times \mathcal{Y}} (\mathbb{1}_{\delta(x, x') \leq r} - \mathbb{1}_{\gamma(y, y') \leq r}) d\pi(x', y') \right| \wedge h dr d\pi(x, y) \leq \\
&\frac{1}{h} \int_{\mathcal{X} \times \mathcal{Y}} \int_{\mathcal{X} \times \mathcal{Y}} \int_0^{+\infty} |\mathbb{1}_{\delta(x, x') \leq r} - \mathbb{1}_{\gamma(y, y') \leq r}| dr d\pi(x', y') d\pi(x, y) = \\
&\frac{1}{h} \int_{\mathcal{X} \times \mathcal{Y}} \int_{\mathcal{X} \times \mathcal{Y}} |\delta(x, x') - \gamma(y, y')| d\pi(x', y') d\pi(x, y),
\end{aligned}$$

which concludes.

Proof of Proposition 2.13

First notice that:

$$\begin{aligned}
W_1(d_{P,h}(P), d_{Q,h}(P)) &\leq \int_{\mathcal{X}} |d_{P,h}(x) - d_{Q,h}(x)| dP(x) \\
&\leq \|d_{P,h} - d_{Q,h}\|_{\infty, \text{Supp}(P)}.
\end{aligned}$$

Then, for all π in $\Pi(P, Q)$:

$$W_1(d_{Q,h}(P), d_{Q,h}(Q)) \leq \int_{\mathcal{X} \times \mathcal{Y}} |d_{Q,h}(x) - d_{Q,h}(y)| d\pi(x, y).$$

Thus, since $d_{Q,h}$ is 1-Lipschitz:

$$W_1(d_{Q,h}(P), d_{Q,h}(Q)) \leq W_1(P, Q).$$

Then, the result follows from Proposition 2.2.

Proof of Proposition 2.15

First notice that $F_{d_{Q,h}(Q)}^{-1} = \lambda F_{d_{P,h}(P)}^{-1}$. Then,

$$\begin{aligned} W_1(d_{P,h}(P), d_{Q,h}(Q)) &= \int_0^1 \left| F_{d_{P,h}(P)}^{-1}(s) - F_{d_{Q,h}(Q)}^{-1}(s) \right| ds \\ &= |1 - \lambda| \int_0^1 \left| F_{d_{P,h}(P)}^{-1}(s) \right| ds \\ &= |1 - \lambda| \mathbb{E}_P [d_{P,h}(X)]. \end{aligned}$$

Proof of Proposition 2.16

This proposition comes from the fact that for h small enough, for both mm-space, the distance-to-measure function will attain its minimum on a set of positive measure. Moreover, the two minima are different since the Lebesgue measure of the open sets are different.

More precisely, if the set $O_{\epsilon(h,O)}$ is non-empty, then the minimal value of the distance-to-measure is given by

$$\min_{x \in \mathbb{R}^d} (d_{P_O,h}(x)) = d_{\min} := \frac{d}{d+1} \left(\frac{h \text{Leb}_d(O)}{\omega_d} \right)^{\frac{1}{d}}.$$

Moreover, the points at minimal distance are exactly the points of $O_{\epsilon(h,O)}$. This is Proposition 2.6. So, $F_{d_{P_O,h}(P_O)}(d_{\min}) = P_O(O_{\epsilon(h,O)})$. To conclude, we use the definition of the L_1 -Wasserstein distance as the L_1 -norm between the cumulative distribution functions.

Proof of Example 2.17

Since $\text{Leb}_d(O) = \omega_d$, $\epsilon(h, O) = h^{\frac{1}{d}}$, and $O_{\epsilon(h,O)}$ is the ball of radius $1 - \epsilon(h, O)$, thus, $P_O(O_{\epsilon(h,O)}) = \left(1 - h^{\frac{1}{d}}\right)^d$.

Also, $\text{Leb}_d(O') = 2^d$, $\epsilon(h, O') = 2 \left(\frac{h}{\omega_d}\right)^{\frac{1}{d}}$, and $O'_{\epsilon(h,O')}$ is the hypercube of radius $1 - \epsilon(h, O')$, thus, $P_{O'}(O'_{\epsilon(h,O')}) = \left(1 - 2 \left(\frac{h}{\omega_d}\right)^{\frac{1}{d}}\right)^d$.

Thus, when h is smaller than $\frac{\omega_d}{2^d}$, the DTM-signature discriminates between P_O and $P_{O'}$. Moreover, the L_1 -Wasserstein distance between the signatures is bounded above by

$$\left(1 - 2 \left(\frac{h}{\omega_d}\right)^{\frac{1}{d}}\right)^d \frac{d}{d+1} \left(\frac{h}{\omega_d}\right)^{\frac{1}{d}} \left(2 - \omega_d^{\frac{1}{d}}\right).$$

Proof of Proposition 2.18

As for Proposition 2.6, we get that for any point x in O :

$$d_{P_O,h}(x) \geq d_{\min} := \left(\frac{h \text{Leb}_d(O)}{\omega_d}\right)^{\frac{1}{d}} \frac{d}{1+d}.$$

We will lower bound the L_1 -Wasserstein distance between $d_{P_O,h}(P_O)$ and $d_{Q,h}(Q)$ by the integral of $F_{d_{Q,h}(Q)}$ over the interval $[0, d_{\min}]$, since $F_{d_{P_O,h}(P_O)}$ equals zero on this interval. We thus need to lower bound $F_{d_{Q,h}(Q)}(t)$ for all $t \leq d_{\min}$.

As for Proposition 2.6, for $\lambda \geq 1$, any point x of $\{f \geq \lambda\}_{\lambda^{-\frac{1}{d}} \left(\frac{h\text{Leb}_d(O)}{\omega_d}\right)^{\frac{1}{d}}}$ satisfies $d_{Q,h}(x) \leq \frac{d_{\min}}{\lambda^{\frac{1}{d}}}$. Thus,

$$F_{d_{Q,h}(Q)} \left(\frac{d_{\min}}{\lambda^{\frac{1}{d}}} \right) \geq Q \left(\{f \geq \lambda\}_{\lambda^{-\frac{1}{d}} \left(\frac{h\text{Leb}_d(O)}{\omega_d}\right)^{\frac{1}{d}}} \right).$$

And we get by denoting $\lambda(t)$ the real number λ satisfying $t = \frac{d_{\min}}{\lambda^{\frac{1}{d}}}$, that:

$$W_1(d_{P_{O,h}(P_O)}, d_{Q,h}(Q)) \geq \int_{t=0}^{d_{\min}} Q \left(\{f \geq \lambda(t)\}_{\lambda(t)^{-\frac{1}{d}} \left(\frac{h\text{Leb}_d(O)}{\omega_d}\right)^{\frac{1}{d}}} \right) dt.$$

Since a cumulative distribution function is non decreasing, we get:

$$\begin{aligned} W_1(d_{P_{O,h}(P_O)}, d_{Q,h}(Q)) &\geq \\ &\int_{t=0}^{d_{\min}} \sup_{t' \leq t} Q \left(\{f \geq \lambda(t')\}_{\lambda(t')^{-\frac{1}{d}} \left(\frac{h\text{Leb}_d(O)}{\omega_d}\right)^{\frac{1}{d}}} \right) dt \\ &= \int_{\lambda=1}^{\infty} d_{\min} \frac{1}{d} \frac{1}{\lambda^{\frac{1}{d}}} \frac{1}{\lambda} \sup_{\lambda' \geq \lambda} Q \left(\{f \geq \lambda'\}_{\lambda'^{-\frac{1}{d}} \left(\frac{h\text{Leb}_d(O)}{\omega_d}\right)^{\frac{1}{d}}} \right) d\lambda \\ &\geq \frac{1}{d+1} \left(\frac{h\text{Leb}_d(O)}{\omega_d} \right)^{\frac{1}{d}} \int_{\lambda=1}^{\infty} \frac{1}{\lambda^{\frac{1}{d}}} \sup_{\lambda' \geq \lambda} P_O \left(\{f \geq \lambda'\}_{\left(\frac{h\text{Leb}_d(O)}{\lambda' \omega_d}\right)^{\frac{1}{d}}} \right) d\lambda. \end{aligned}$$

■

Now we assume that the density f is Hölder over O with parameters χ in $[0, 1]$ and L in \mathbb{R}_+^* .

Proof of Proposition 2.19

First notice that for all positive λ , with $\epsilon(\lambda) = \lambda^{-\frac{1}{d}} \left(\frac{h\text{Leb}_d(O)}{\omega_d}\right)^{\frac{1}{d}}$ we have:

$$\{f \geq \lambda + L\epsilon(\lambda)^\chi\} \cap O_{\epsilon(\lambda)} \subset \{f \geq \lambda\}_{\epsilon(\lambda)}.$$

According to Proposition 2.18, the aim is thus to show that for some λ larger than 1, the set $\{f \geq \lambda + L\epsilon(\lambda)^\chi\} \cap O_{\epsilon(\lambda)}$ is non-empty. We thus focus on the supremum of f over $O_{\epsilon(\lambda)}$, which we denote by $\|f\|_{\infty, \epsilon(\lambda)}$ and try to prove that it is larger than $\lambda + L\epsilon(\lambda)^\chi$.

Remind that if $\text{Reach}(O) \geq \epsilon(\lambda)$, then according to Proposition 2.7, the set $O_{\epsilon(\lambda)}^{\epsilon(\lambda)}$ equals O . Since f is Hölder, we can thus build some sequence $(y_n)_{n \in \mathbb{N}^*}$ in $O_{\epsilon(\lambda)}$, such that $f(y_n) \geq \|f\|_{\infty, O} - \frac{1}{n} - L\epsilon(\lambda)^\chi$. Finally we get:

$$\|f\|_{\infty, \epsilon(\lambda)} \geq \|f\|_{\infty, O} - L\epsilon(\lambda)^\chi.$$

So the quantity $W_1(d_{P_{O,h}(P_O)}, d_{Q,h}(Q))$ is positive whenever:

$$\|f\|_{\infty, O} > \inf \{ \lambda + 2L\epsilon(\lambda)^\chi \mid \lambda \geq 1, \epsilon(\lambda) \leq \text{Reach}(O) \}.$$

With $\lambda_0 = 1$, we have $\lambda_0 + 2L\epsilon(\lambda_0)^\chi = 1 + 2L \left(\frac{h\text{Leb}_d(O)}{\omega_d}\right)^{\frac{\chi}{d}}$.

With λ_1 satisfying $\epsilon(\lambda_1) = \text{Reach}(O)$, we have:

$$\lambda_1 + 2L\epsilon(\lambda_1)^\chi = \frac{1}{(\text{Reach}(O))^d} \frac{h\text{Leb}_d(O)}{\omega_d} + 2L(\text{Reach}(O)^\chi).$$

We also have that

$$\inf \{ \lambda + 2L\epsilon(\lambda)^\chi \mid \lambda > 0 \} = (2L)^{\frac{d}{d+\chi}} \left(\frac{\text{Leb}_d(O)}{\omega_d} \right)^{\frac{\chi}{d+\chi}} h^{\frac{\chi}{d+\chi}} \left[\left(\frac{\chi}{d} \right)^{\frac{d}{d+\chi}} + \left(\frac{\chi}{d} \right)^{-\frac{\chi}{d+\chi}} \right].$$

The infimum is attained at $\lambda_2 = \left(\frac{\chi}{d} \right)^{\frac{d}{d+\chi}} (2L)^{\frac{d}{d+\chi}} \left(\frac{h\text{Leb}_d(O)}{\omega_d} \right)^{\frac{\chi}{d+\chi}}$.

It proves the first part of the proposition.

The second part is a straightforward consequence of the proof of Proposition 2.18.

Proof of Example 2.20

The measure Q is absolutely continuous with respect to P_O with density f defined by

$$f(x) = C_\sigma \frac{1}{(2\pi)^{\frac{d}{2}} \sigma^d} \exp -\frac{\|x\|_2^2}{2\sigma^2},$$

with

$$C_\sigma = \frac{1}{\int_{\mathcal{B}(0,1)} \frac{1}{(2\pi)^{\frac{d}{2}} \sigma^d} \exp -\frac{\|x\|_2^2}{2\sigma^2} dx}.$$

According to the previous proposition, we can prove that the signatures are discriminative provided that σ is small enough, for all h smaller than 0.23 if $d = 1$; 0.30 if $d = 2$; 0.68 if $d = 3$; and for all value of h is $d \geq 4$.

More precisely, the signatures are discriminative when

$$h \leq \sigma^d \frac{\exp \frac{d}{2}}{2^d} \left(1 - \frac{(2\pi)^{\frac{d}{2}} \sigma^d}{C_\sigma} \right)^d,$$

or when

$$h \in \left[\sigma^{d^2+d} \left(\frac{d}{2} \right)^d \frac{(2\pi)^{\frac{d^2}{2}}}{C_\sigma^d} \exp \frac{d}{2}, \min \left\{ 1, C_\sigma \exp \frac{d}{2} \left(\frac{d}{2} \right)^d \left(\frac{1}{d+1} \right)^{\frac{1}{d+1}}, \frac{2}{d} C_\sigma \frac{\exp -\frac{1}{2}}{(2\pi)^{\frac{d}{2}} \sigma^{d+1}} \right\} \right).$$

When σ is small enough, this last assumption can be rewritten as follows

$$h \in \left[\sigma^{d^2+d} \left(\frac{d}{2} \right)^d \frac{(2\pi)^{\frac{d^2}{2}}}{C_\sigma^d} \exp \frac{d}{2}, \min \{ 1, C_\sigma C'_d \} \right),$$

with $C'_d = \frac{2}{d} \frac{\exp -\frac{1}{2}}{(2\pi)^{\frac{d}{2}} \sigma^{d+1}}$. Note that C'_d is much bigger than 1 when $d \geq 4$. Also, $C'_1 \simeq 0.23$, $C'_2 \simeq 0.30$, $C'_3 \simeq 0.68$.

In order to get these results, we use Proposition 2.19. According to the mean value theorem on $\mathcal{B}(0, 1)$, which is a convex subset of \mathbb{R}^d , the density f is Lipschitz with parameter

$$L = C_\sigma \frac{\exp -\frac{1}{2}}{(2\pi)^{\frac{d}{2}} \sigma^{d+1}}.$$

Thus we use $\chi = 1$ since f is Lipschitz. Moreover, the reach of a ball is equal to its radius, thus $\text{Reach}(O) = 1$, and $\text{Leb}_d(O) = \omega_d$ by definition.

Note that when σ goes to zero, the scaling parameter C_σ goes to 1.

Proof of Lemma 2.4

We give a proof for Lemma 2.4 that was stated by Lieutier in [Lie04]. I am extremely grateful to Frédéric Chazal for his help.

First, we need to introduce some setting and recall some results from [Lie04].

For $x \in \mathbb{R}^d$, we set $\Gamma(x)$ the set of all projections of x on ∂O :

$$\Gamma(x) = \{y \in \partial O \mid \|x - y\|_2 = d_{\partial O}(x)\}.$$

Also, we denote by $\mathcal{R}(x) = d_{\partial O}(x)$ the distance of x to the boundary ∂O .

Definition 2.32. We define the function \mathcal{F} on O by:

$$\mathcal{F}(x) = \inf\{r \geq 0 \mid \exists y \in \mathbb{R}^d, \overline{\mathcal{B}}(y, r) \supset \Gamma(x)\}.$$

Note that this function is upper bounded by the diameter of the bounded open set O . Moreover,

$$\mathcal{F}(x) > 0 \Leftrightarrow x \in \mathcal{M}(O).$$

Definition 2.33. We defined by $\Theta(x)$ the centre of the closed ball with radius $\mathcal{F}(x)$ that is minimal among all closed balls that contain $\Gamma(x)$.

Such a minimal ball exists and is unique. Indeed, if not, a ball containing the union of two balls containing $\Gamma(x)$ would also contain $\Gamma(x)$ and would have a smaller radius.




Figure 2.22: Γ , \mathcal{R} , \mathcal{F} and Θ . Figure from [CL05]

The vectorial function ∇ that generalises the gradient function is defined as follows:

Definition 2.34. The vectorial function ∇ is defined by:

$$\nabla(x) = \frac{x - \Theta(x)}{\mathcal{R}(x)}.$$

According to [Lie04], the following proposition is satisfied:

Proposition 2.35 ([Lie04]). *If O is a bounded open set, it comes:*

$$\|\nabla(x)\|^2 + \frac{\mathcal{F}(x)^2}{\mathcal{R}(x)^2} = 1.$$

As a consequence,

$$0 \leq \|\nabla(x)\| \leq 1 \text{ and } \|\nabla(x)\| < 1 \Leftrightarrow x \in \mathcal{M}(O).$$

Lieutier introduces in [Lie04] a continuous map $C : \mathbb{R}_+ \times O \rightarrow O$ that satisfies for every $x \in O$, and $t \in \left[0, \frac{\mathcal{R}(x)}{4}\right]$:

$$C(t, x) = x + \int_0^t \nabla(C(\tau, x)) d\tau,$$

and

$$\mathcal{R}(C(t, x)) = \mathcal{R}(x) + \int_0^t \|\nabla(C(\tau, x))\|^2 d\tau.$$

Moreover, the image of the curve $t \mapsto C(t, x)$ is a rectifiable curve. Its arc length is an increasing function of t defined by:

$$s(t) = \int_0^t \|\nabla(C(\tau, x))\| d\tau.$$

As a consequence, the map $t \mapsto \mathcal{R}(C(t, x))$ is non-decreasing and 1-Lipschitz.

Lieutier also proves that the map $t \mapsto \mathcal{F}(C(t, x))$ is increasing. So, $H : [0, 1] \times O \rightarrow O$, $(t, x) \mapsto C(tD, x)$ is a retraction of O onto its medial axis, where D denotes the diameter of O .

Finally, Lieutier proves the following Proposition:

Proposition 2.36 ([Lie04]). *Given $T > 0$, if $\nabla(C(T, x)) \neq 0$, then the arc-length map $s : [0, T] \rightarrow [0, s(T)]$, $t \mapsto s(t)$ is increasing, and has an inverse function defined by: $t : [0, s(T)] \rightarrow [0, T]$, $s \mapsto t(s)$. As a consequence, it is possible to re-parametrise the curve with the arc-length map, and it comes that:*

$$\mathcal{R}(C(t(s), x)) = \mathcal{R}(x) + \int_0^s \|\nabla(C(t(\sigma), x))\| d\sigma.$$

Now we prove Lemma 2.4.

First we prove that the skeleton is a subset of the closure of the medial axis. Set x in the skeleton. For the sake of contradiction, assume that x is not in the closure of the medial axis. Set $p(x)$ the unique element of $\Gamma(x)$. Then a ball \mathcal{B}_0 centred at x is included in $\mathcal{M}(O)^c \cap O$. Thus, the gradient $\nabla(x)$ is non-zero on this neighbourhood \mathcal{B}_0 . The continuity of the map $s \mapsto C(t(s), x)$ entails that for $T > 0$ small enough, $\nabla(C(T, x)) \neq 0$. Thus, for every $s \leq s(T)$, $C(t(s), x)$ is in \mathcal{B}_0 . According to Proposition 2.36, $\mathcal{R}(C(t(s), x)) = \mathcal{R}(x) + \int_0^s \|\nabla(C(t(\sigma), x))\| d\sigma$, with $\nabla(C(t(\sigma), x)) = 1 \forall \sigma \leq s$. In particular, we deduce that:

$$d_{\partial O}(C(T, x)) = d_{\partial O}(x) + s.$$

So,

$$\|C(T, x) - p(x)\| \geq \|x - p(x)\| + s.$$

Moreover, since we parametrised the curve with arc-length,

$$\|C(T, x) - x\| \leq s.$$

As a consequence, equality holds everywhere in the following sequence of inequalities:

$$\begin{aligned} \|C(T, x) - p(x)\| &\leq \|C(T, x) - x\| + \|x - p(x)\| \\ &\leq \|C(T, x) - x\| + \|C(T, x) - p(x)\| - s \\ &\leq \|C(T, x) - p(x)\|. \end{aligned}$$

So, $\|C(T, x) - x\| = s$ and $C(T, x)$ lies on the half line $(x, p(x)] \setminus]x, p(x)]$. As a consequence, the ball centred at $C(T, x)$ with radius $s + \|x - p(x)\|$ contains the ball $\mathcal{B}(x, \mathcal{R}(x))$ and is included in O' , since $\mathcal{R}(C(T, x)) = s + \|x - p(x)\|$, with $s > 0$. This is a contradiction with the maximality of the ball $\mathcal{B}(x, \mathcal{R}(x))$ and to the fact that x was in the skeleton.

The inclusion of the medial axis into the skeleton is easier. It comes from the fact that for any point x in the medial axis, the ball \mathcal{B} centred at x with radius $\mathcal{R}(x)$ is maximal for the inclusion order. Indeed, any ball containing \mathcal{B} and included in O should be such that $\partial O \cap \partial \mathcal{B}(x, \mathcal{R}(x))$ contains the projection of x on ∂O , thus should contain at least two elements. These balls should then coincide. As a consequence \mathcal{B} is maximal and x is in the skeleton, hence the inclusion.

2.3.2 Proofs for Section 2.2

Proof of Lemma 2.21

Note that for $(X'_1, X'_2, \dots, X'_N)$ an N -sample of law P and ϕ an isomorphism between (\mathcal{X}, δ, P) and (\mathcal{Y}, γ, Q) , the tuple $(\phi(X'_1), \phi(X'_2), \dots, \phi(X'_N))$ is an N -sample of law Q . Moreover, $\delta(X'_i, X'_j) = \gamma(\phi(X'_i), \phi(X'_j))$ for all i and j in $\llbracket 1, N \rrbracket$. It follows that the distances and the nearest neighbours are preserved.

Thus, the distributions of $(d_{\hat{P}_N, h}(X'_i))_{i \in \llbracket 1, n \rrbracket}$ and $(d_{\hat{Q}_N, h}(Y_i))_{i \in \llbracket 1, n \rrbracket}$ are equal.

The lemma follows from the equality:

$$\begin{aligned} & W_1(d_{\hat{P}_N, h}(\hat{P}_n), d_{\hat{Q}_N, h}(\hat{Q}_n)) \\ &= \int_0^{+\infty} \frac{1}{n} \left| \sum_{i=1}^n \mathbb{1}_{d_{\hat{P}_N, h}(X_i) \leq s} - \sum_{i=1}^n \mathbb{1}_{d_{\hat{Q}_N, h}(Y_i) \leq s} \right| ds, \end{aligned}$$

with (X_1, X_2, \dots, X_N) an N -sample from P .

L_1 -Wasserstein distance between the distributions of interest

Lemma 2.37. *The quantity $W_1(\mathcal{L}_{N, n, h}(P, P), \mathcal{L}_{N, n, h}^*(\hat{P}_N, \hat{P}_N))$ is bounded above by*

$$2\sqrt{n} \left(\mathbb{E}[\|d_{\hat{P}_N, h} - d_{P, h}\|_{\infty, \mathcal{X}}] + W_1(d_{P, h}(P), d_{P, h}(\hat{P}_N)) + \|d_{P, h} - d_{\hat{P}_N, h}\|_{\infty, \mathcal{X}} \right).$$

Proof of Lemma 2.37. Let (X_1, X_2, \dots, X_N) be an N -sample of law P , and \hat{P}_N the associated empirical measure. We can upper bound the L_1 -Wasserstein distance between the subsampling distribution

$$\mathcal{L}^*(\sqrt{n}W_1(d_{\hat{P}_N, h}(P_n^*), d_{\hat{P}_N, h}(P_n^*)) | \hat{P}_N)$$

and the distribution of interest

$$\mathcal{L}(\sqrt{n}W_1(d_{\hat{P}_N, h}(\hat{P}_n), d_{\hat{P}'_N, h}(\hat{P}'_n))),$$

by

$$W_1 \left(\mathcal{L} \left(\sqrt{n}W_1 \left(d_{\hat{P}_N, h}(P_n^*), d_{\hat{P}_N, h}(P_n^*) \right) | \hat{P}_N \right), \mathcal{L} \left(\sqrt{n}W_1 \left(d_{P, h}(P_n^*), d_{P, h}(P_n^*) \right) | \hat{P}_N \right) \right) \quad (2.4)$$

$$+ W_1 \left(\mathcal{L} \left(\sqrt{n}W_1 \left(d_{P, h}(P_n^*), d_{P, h}(P_n^*) \right) | \hat{P}_N \right), \mathcal{L} \left(\sqrt{n}W_1 \left(d_{P, h}(\hat{P}_n), d_{P, h}(\hat{P}_n) \right) \right) \right) \quad (2.5)$$

$$+ W_1 \left(\mathcal{L} \left(\sqrt{n}W_1 \left(d_{P, h}(\hat{P}_n), d_{P, h}(\hat{P}_n) \right) \right), \mathcal{L} \left(\sqrt{n}W_1 \left(d_{\hat{P}_N, h}(\hat{P}_n), d_{\hat{P}'_N, h}(\hat{P}'_n) \right) \right) \right). \quad (2.6)$$

We bound the term 2.4 by

$$2\sqrt{n} \|d_{P, h} - d_{\hat{P}_N, h}\|_{\infty, \mathcal{X}}.$$

the term 2.5 by

$$2\sqrt{n}W_1\left(d_{P,h}(P), d_{P,h}(\hat{P}_N)\right)$$

and the term 2.6 by

$$2\sqrt{n}\mathbb{E}[\|d_{P,h} - d_{\hat{P}_N,h}\|_{\infty,\mathcal{X}}].$$

This is proved in the three following lemmas. □

Lemma 2.38 (Study of term 2.6). *We have*

$$W_1\left(\mathcal{L}(\sqrt{n}W_1(d_{\hat{P}_N,h}(\hat{P}_n), d_{\hat{P}'_N,h}(\hat{P}'_n))), \mathcal{L}(\sqrt{n}W_1(d_{P,h}(\hat{P}_n), d_{P,h}(\hat{P}'_n)))\right) \leq 2\sqrt{n}\mathbb{E}[\|d_{P,h} - d_{\hat{P}_N,h}\|_{\infty,\mathcal{X}}].$$

Proof of Lemma 2.38. To bound this L_1 -Wasserstein distance, we choose as a transport plan the law of the random vector

$$(\sqrt{n}W_1(d_{\hat{P}_N,h}(\hat{P}_n), d_{\hat{P}'_N,h}(\hat{P}'_n)), \sqrt{n}W_1(d_{P,h}(\hat{P}_n), d_{P,h}(\hat{P}'_n))),$$

with $\hat{P}_n, \hat{P}'_n, \hat{P}_{N-n}$ and \hat{P}'_{N-n} independent empirical measures of law P . Then the L_1 -Wasserstein distance is bounded by

$$\mathbb{E}[|\sqrt{n}W_1(d_{\hat{P}_N,h}(\hat{P}_n), d_{\hat{P}'_N,h}(\hat{P}'_n)) - \sqrt{n}W_1(d_{P,h}(\hat{P}_n), d_{P,h}(\hat{P}'_n))|],$$

which is not bigger than:

$$\sqrt{n}\mathbb{E}[W_1(d_{\hat{P}_N,h}(\hat{P}_n), d_{P,h}(\hat{P}_n)) + W_1(d_{\hat{P}'_N,h}(\hat{P}'_n), d_{P,h}(\hat{P}'_n))].$$

We bound the term $\mathbb{E}[W_1(d_{\hat{P}_N,h}(\hat{P}_n), d_{P,h}(\hat{P}_n))]$ by $\mathbb{E}[\|d_{P,h} - d_{\hat{P}_N,h}\|_{\infty,\mathcal{X}}]$, thanks to Lemma 2.41. □

Lemma 2.39 (Study of term 2.5). *We have*

$$W_1\left(\mathcal{L}(\sqrt{n}W_1(d_{P,h}(\hat{P}_n), d_{P,h}(\hat{P}'_n)), \mathcal{L}(\sqrt{n}W_1(d_{P,h}(P_n^*), d_{P,h}(P_n'^*))|\hat{P}_N)\right) \leq 2\sqrt{n}W_1(d_{P,h}(P), d_{P,h}(\hat{P}_N)).$$

Proof of Lemma 2.39. Let π be the optimal transport plan associated to $W_1(d_{P,h}(P), d_{P,h}(\hat{P}_N))$; see the definition of the L_1 -Wasserstein with transport plans.

From an n -sample of law π , we get two empirical distributions $d_{P,h}(\hat{P}_n)$ and $d_{P,h}(P_n^*)$. Independently, from another n -sample of law π , we get $d_{P,h}(\hat{P}'_n)$ and $d_{P,h}(P_n'^*)$.

The L_1 -Wasserstein distance is then bounded by

$$\sqrt{n}\mathbb{E}_{\pi^{\otimes n} \otimes \pi^{\otimes n}}[W_1(d_{P,h}(\hat{P}_n), d_{P,h}(P_n^*)) + W_1(d_{P,h}(\hat{P}'_n), d_{P,h}(P_n'^*))].$$

Now notice that, if we denote $\hat{P}_n = \sum_{i=1}^n \frac{1}{n} \delta_{Y_i}$ and $P_n^* = \sum_{i=1}^n \frac{1}{n} \delta_{Z_i}$, we have:

$$\begin{aligned} W_1(d_{P,h}(\hat{P}_n), d_{P,h}(P_n^*)) &= \int_{t=0}^{+\infty} \left| \frac{1}{n} \sum_{i=1}^n \mathbb{1}_{d_{P,h}(Y_i) \leq t} - \frac{1}{n} \sum_{i=1}^n \mathbb{1}_{d_{P,h}(Z_i) \leq t} \right| dt \\ &\leq \frac{1}{n} \sum_{i=1}^n \int_{t=0}^{+\infty} \left| \mathbb{1}_{d_{P,h}(Y_i) \leq t} - \mathbb{1}_{d_{P,h}(Z_i) \leq t} \right| dt \\ &= \frac{1}{n} \sum_{i=1}^n |d_{P,h}(Y_i) - d_{P,h}(Z_i)|. \end{aligned}$$

So, the L_1 -Wasserstein distance is not bigger than

$$2\sqrt{n}\mathbb{E}[|d_{P,h}(Y) - d_{P,h}(Z)|],$$

with $(d_{P,h}(Y), d_{P,h}(Z))$ of law π , so we get the upper bound:

$$2\sqrt{n} \left(W_1(d_{P,h}(P), d_{P,h}(\hat{P}_N)) \right).$$

□

Lemma 2.40 (Study of term 2.4). *We have*

$$W_1 \left(\mathcal{L}(\sqrt{n}W_1(d_{P,h}(P_n^*), d_{P,h}(P_n^*))|\hat{P}_N), \mathcal{L}(\sqrt{n}W_1(d_{\hat{P}_N,h}(P_n^*), d_{\hat{P}_N,h}(P_n^*))|\hat{P}_N) \right) \leq 2\sqrt{n}\|d_{P,h} - d_{\hat{P}_N,h}\|_{\infty, \mathcal{X}}.$$

Proof of Lemma 2.4. It is the same proof as for the first lemma, except that \hat{P}_N is fixed. □

Lemma 2.41. *Let Q, P and P' be some measures over some metric space (\mathcal{X}, δ) , we have:*

$$W_1(d_{P,h}(Q), d_{P',h}(Q)) \leq \int_{\mathcal{X}} |d_{P,h}(x) - d_{P',h}(x)| dQ(x) \leq \|d_{P,h} - d_{P',h}\|_{\infty, \text{Supp}(Q)}.$$

Proof of Lemma 2.41. We chose the transport plan $(d_{P,h}(Y), d_{P',h}(Y))$ for Y of distribution Q . □

Thanks to Proposition 2.2 and to the fact that the distance-to-measure is 1-Lipschitz, we can derive another upper bound depending only on the L_1 -Wasserstein distance between the measure P and its empirical versions:

Corollary 2.42. *The quantity $W_1 \left(\mathcal{L}_{N,n,h}(P, P), \mathcal{L}_{N,n,h}^*(\hat{P}_N, \hat{P}_N) \right)$ is bounded above by*

$$2\frac{\sqrt{n}}{h}\mathbb{E} \left[W_1 \left(\hat{P}_N, P \right) \right] + 2\sqrt{n} \left(1 + \frac{1}{h} \right) W_1(\hat{P}_N, P).$$

The rates of convergence of the L_1 -Wasserstein distance between a Borel probability measure on the Euclidean space \mathbb{R}^d and its empirical version are faster when the dimension d is low; see [FG15]. Thus, we prefer to use the first bound for regular measures. In this case, we use rates of convergence for the distance-to-measure, derived in [CMM16]. For regular measures, in some cases, the bound in Lemma 2.37 is better than the bound in Corollary 2.42.

Proof of Lemma 2.24

The random function $\sqrt{n} \left(F_{d_{P,h}(P)} - F_{d_{P,h}(\hat{P}_n)} \right)$ converges weakly in L_1 to some gaussian process $\mathbb{G}_{P,h}$ with covariance kernel

$$\kappa(s, t) = F_{d_{P,h}(P)}(s) \left(1 - F_{d_{P,h}(P)}(t) \right)$$

for $s \leq t$; see [dBG99] or part 3.3 of [BL14]. Thanks to Theorem 2.8, p.23, in [Bil99], since $L_1 \times L_1$ is separable and \hat{P}_n and \hat{P}'_n are independent, the random vector

$$\left(\sqrt{n} \left(F_{d_{P,h}(P)} - F_{d_{P,h}(\hat{P}_n)} \right), \sqrt{n} \left(F_{d_{P,h}(P)} - F_{d_{P,h}(\hat{P}'_n)} \right) \right)$$

converges weakly to $(\mathbb{G}_{P,h}, \mathbb{G}'_{P,h})$ with $\mathbb{G}_{P,h}$ and $\mathbb{G}'_{P,h}$ independent Gaussian processes. Since the map $(x, y) \mapsto x - y$ is continuous in L_1 , the mapping theorem states that $\sqrt{n} \left(F_{d_{P,h}(\hat{P}'_n)} - F_{d_{P,h}(\hat{P}_n)} \right)$

converges weakly to the Gaussian process $\mathbb{G}_{P,h} - \mathbb{G}'_{P,h}$ in L_1 . Once more we use the mapping theorem with the continuous map $x \mapsto \|x\|_1$ and the definition of the L_1 -Wasserstein distance as the L_1 -norm of the cumulative distribution functions to get that:

$$\sqrt{n}W_1(d_{P,h}(\hat{P}_n), d_{P,h}(\hat{P}'_n)) \rightsquigarrow \|\mathbb{G}_{P,h} - \mathbb{G}'_{P,h}\|_1.$$

We then get the convergence of moments following the same method as for Theorem 2.4 in [dBGM99]. We have the bound $\mathbb{E}[\|t \mapsto \mathbb{1}_{d_{P,h}(X_i) \leq t} - \mathbb{1}_{d_{P,h}(Y_i) \leq t}\|_1] \leq \mathcal{D}_P < \infty$. Moreover, the random function $\sqrt{n} \left(F_{d_{P,h}(\hat{P}'_n)} - F_{d_{P,h}(\hat{P}_n)} \right)$ converges weakly to the gaussian process $\mathbb{G}_{P,h} - \mathbb{G}'_{P,h}$ in L_1 . So, thanks to Theorem 5.1 in [dAG79] (cited in [AG80] p.136), we have:

$$\mathbb{E}[\sqrt{n}W_1(d_{P,h}(\hat{P}_n), d_{P,h}(\hat{P}'_n))] \rightarrow \mathbb{E}[\|\mathbb{G}_{P,h} - \mathbb{G}'_{P,h}\|_1].$$

We deduce that:

$$W_1 \left(\mathcal{L} \left(\sqrt{n}W_1 \left(d_{P,h}(\hat{P}_n), d_{P,h}(\hat{P}'_n) \right) \right), \mathcal{L} \left(\|\mathbb{G}_{P,h} - \mathbb{G}'_{P,h}\|_1 \right) \right) \rightarrow 0.$$

Moreover, we have the bound:

$$W_1 \left(\mathcal{L} \left(\sqrt{n}W_1 \left(d_{P,h}(\hat{P}_n), d_{P,h}(\hat{P}'_n) \right) \right), \mathcal{L}_{N,n,h}(P, P) \right) \leq 2\sqrt{n}\mathbb{E}[\|d_{P,h} - d_{\hat{P}_N,h}\|_{\infty, \mathcal{X}}].$$

So, if $\sqrt{n}\mathbb{E}[\|d_{P,h} - d_{\hat{P}_N,h}\|_{\infty, \mathcal{X}}] \rightarrow 0$ when $N \rightarrow \infty$, we have that:

$$W_1 \left(\mathcal{L}_{N,n,h}(P, P), \mathcal{L} \left(\|\mathbb{G}_{P,h} - \mathbb{G}'_{P,h}\|_1 \right) \right) \rightarrow 0.$$

Finally, with the same arguments as for Lemma 2.37, we get that:

$$\begin{aligned} & W_1 \left(\mathcal{L}_{N,n,h}^*(\hat{P}_N, \hat{P}_N), \mathcal{L} \left(\|\mathbb{G}_{P,h} - \mathbb{G}'_{P,h}\|_1 \right) \right) \leq \\ & W_1 \left(\mathcal{L} \left(\sqrt{n}W_1 \left(d_{P,h}(\hat{P}_n), d_{P,h}(\hat{P}'_n) \right) \right), \mathcal{L} \left(\|\mathbb{G}_{P,h} - \mathbb{G}'_{P,h}\|_1 \right) \right) \\ & + 2\sqrt{n}W_1 \left(d_{P,h}(P), d_{P,h}(\hat{P}_N) \right) + 2\sqrt{n}\|d_{P,h} - d_{\hat{P}_N,h}\|_{\infty, \mathcal{X}}. \end{aligned}$$

Proof of Lemma 2.25

Let $\epsilon < \alpha$ and η be two positive numbers.

The probability $\mathbb{P}_{(P,Q)}(\phi_{N,n,h} = 1)$ is bounded above by

$$\mathbb{P} \left(\sqrt{n}W_1 \left(d_{\hat{P}_N,h}(\hat{P}_n), d_{\hat{Q}_N,h}(\hat{Q}_n) \right) \geq q_{\alpha+\epsilon} - \eta \right) + \mathbb{P}(\hat{q}_\alpha < q_{\alpha+\epsilon} - \eta).$$

With a drawing, we see that $\mathbb{P}(\hat{q}_\alpha < q_{\alpha+\epsilon} - \eta)$ is bounded above by

$$\mathbb{P} \left(W_1 \left(\mathcal{L} \left(\frac{1}{2}\|\mathbb{G}_{P,h} - \mathbb{G}'_{P,h}\|_1 + \frac{1}{2}\|\mathbb{G}_{Q,h} - \mathbb{G}'_{Q,h}\|_1 \right), \mathcal{L}^* \right) \geq \epsilon\eta \right),$$

where $\mathcal{L}^* = \frac{1}{2}\mathcal{L}_{N,n,h}^*(\hat{P}_N, \hat{P}_N) + \frac{1}{2}\mathcal{L}_{N,n,h}^*(\hat{Q}_N, \hat{Q}_N)$.

Thanks to the weak convergences in Lemma 2.24 of the paper and the Portmanteau lemma, $\limsup_{N \rightarrow \infty} \mathbb{P}_{(P,Q)}(\phi_{N,n,h} = 1)$ is thus bounded above by

$$\mathbb{P} \left(\frac{1}{2}\|\mathbb{G}_{P,h} - \mathbb{G}'_{P,h}\|_1 + \frac{1}{2}\|\mathbb{G}_{Q,h} - \mathbb{G}'_{Q,h}\|_1 \geq q_{\alpha+\epsilon} - \eta \right).$$

We now make η and ϵ go to zero and under the continuity assumption,

$$\limsup_{N \rightarrow \infty} \mathbb{P}_{(P,Q)}(\phi_{N,n,h} = 1) \leq \alpha.$$

As well, we get that $\liminf_{N \rightarrow \infty} \mathbb{P}_{(P,Q)}(\phi_{N,n,h} = 1) \geq \alpha$.

Proof of Proposition 2.26

Proof of part 2 of Proposition 2.26 We may assume that the diameter \mathcal{D}_P of the support of the measure P equals 1. Indeed, if we apply a dilatation to the measure to make the diameter of its support be equal to 1, then the quantity $W_1\left(\mathcal{L}_{N,n,h}(P, P), \mathcal{L}_{N,n,h}^*(\hat{P}_N, \hat{P}_N)\right)$ is simply multiplied by the parameter of the dilatation. By using Corollary 2.42 and Theorem 1 of [FG15], we have a bound for the expectation:

$$\mathbb{E}\left[W_1\left(\mathcal{L}_{N,n,h}(P, P), \mathcal{L}_{N,n,h}^*(\hat{P}_N, \hat{P}_N)\right)\right] \leq \begin{cases} C \frac{\sqrt{n}}{h} N^{-\frac{1}{d}} & \text{if } d > 2 \\ C \frac{\sqrt{n}}{h} N^{-\frac{1}{2}} \log(1+N) & \text{if } d = 2 \\ C \frac{\sqrt{n}}{h} N^{-\frac{1}{2}} & \text{if } d < 2 \end{cases}$$

for some positive constant C depending on P . ■

Proof of part 3 of Proposition 2.26 First notice that for $\lambda > 1$,

$$\mathbb{P}\left(W_1\left(\mathcal{L}_{N,n,h}(P, P), \mathcal{L}_{N,n,h}^*(\hat{P}_N, \hat{P}_N)\right) \geq \lambda\right) = 0$$

under the assumption $\mathcal{D}_P = 1$. We thus focus on values of λ not bigger than 1. In this case, with the Theorem 2 of [FG15], we get easily that:

$$\mathbb{P}\left(W_1\left(\mathcal{L}_{N,n,h}(P, P), \mathcal{L}_{N,n,h}^*(\hat{P}_N, \hat{P}_N)\right) \geq \lambda\right) \leq \begin{cases} C \exp\left(-C' \left(\lambda \frac{N^{\frac{1}{d}} h}{\sqrt{n}} - C''\right)^d\right) & \text{for } d > 2 \\ C \exp\left(-C' \left(\frac{\frac{\sqrt{N} h}{\sqrt{n}} \lambda - C'' \sqrt{\frac{N}{N-n}} \log(1+N-n)}{\log\left(2 + \frac{2\sqrt{N}}{\frac{\sqrt{N} h}{\sqrt{n}} \lambda - C'' \sqrt{\frac{N}{N-n}} \log(1+N-n)}\right)}\right)^2\right) & \text{for } d = 2 \\ C \exp\left(-C' \left(\lambda \frac{\sqrt{N} h}{\sqrt{n}} - C''\right)^2\right) & \text{for } d < 2 \end{cases}$$

for some positive constants C , C' and C'' depending on P .

We conclude the proof with the Borel–Cantelli lemma.

Proof of part 1 of Proposition 2.26 We need to show that under the assumption $\rho > \frac{\max\{d, 2\}}{2}$, the following properties are satisfied:

$$\sqrt{n} \mathbb{E}[\|d_{P,h} - d_{\hat{P}_N,h}\|_{\infty, \mathcal{X}}] \rightarrow 0,$$

$$\sqrt{n} W_1(d_{P,h}(P), d_{P,h}(\hat{P}_N)) \rightarrow 0 \text{ a.e.},$$

and

$$\sqrt{n} \|d_{P,h} - d_{\hat{P}_N,h}\|_{\infty, \mathcal{X}} \rightarrow 0 \text{ a.e.}$$

We treat the case $d > 2$. The cases $d < 2$ and $d = 2$ are similar.

Thanks to Theorem 1 of [FG15], there is some positive constant C depending on P such that for N big enough:

$$\mathbb{E}[W_1(\hat{P}_N, P)] \leq CN^{-\frac{1}{d}}.$$

Thus, thanks to part 2 of Proposition 2.26, the quantity $\sqrt{n}\mathbb{E}[\|d_{P,h} - d_{\hat{P}_N,h}\|_{\infty,\mathcal{X}}]$ goes to zero if $\frac{\sqrt{n}}{h}N^{-\frac{1}{d}}$ goes to zero when N goes to infinity. So, this convergence occurs under the assumption $\rho > \frac{d}{2}$.

We get from Theorem 2 of [FG15] that for $x \leq 1$, there are some positive constants C and c depending on P such that:

$$\mathbb{P}(W_1(\hat{P}_N, P) \geq x) \leq C \exp(-cNx^d).$$

We use this inequality with $x = \frac{h}{\sqrt{n}K}$ for positive integers K . Thanks to the Borel–Cantelli lemma, under the assumption $\rho > \frac{d}{2}$, we get that:

$$\frac{\sqrt{n}}{h}W_1(P, \hat{P}_N) \rightarrow 0 \text{ a.e..}$$

So, thanks to Proposition 2.2, the third property is true.

To finish, note that $d_{P,h}(\hat{P}_N)$ is the empirical measure associated to $d_{P,h}(P)$. Once more we use Theorem 2 of [FG15] and get that for $x \leq 1$, $\mathbb{P}(\sqrt{n}W_1(d_{P,h}(\hat{P}_N), d_{P,h}(P)) \geq x) \leq C \exp(-c\frac{N}{n}x^2)$. Thanks to the Borel–Cantelli lemma, under the assumption $\rho > 1$, the a.e. convergence to zero of $\sqrt{n}W_1(d_{P,h}(P), d_{P,h}(\hat{P}_N))$ occurs.

Proof of Example 2.27

For every point x in O and $r > 0$, according to Lemma 2.31 there exist a maximal ball $\mathcal{B}(x', r')$ included in $O \cap \mathcal{B}(x, r)$ which contains x . Assume for the sake of contradiction that $r' < \min\{\frac{r}{2}, \text{Reach}(O)\}$.

Since $r' < \frac{r}{2}$, the ball $\overline{\mathcal{B}}(x', r')$ is included in $\mathcal{B}(x, r)$ thus $\mathcal{B}(x', r')$ is maximal in O . So x' belongs to $Sk(O)$, and thanks to Lemma 2.4, to $\overline{\mathcal{M}(O)}$. But $r' < \text{Reach}(O)$; this is absurd.

It follows that:

$$P_O(\mathcal{B}(x, r)) \geq P_O\left(\mathcal{B}\left(x', \min\left\{\text{Reach}(O), \frac{r}{2}\right\}\right)\right).$$

So, for $r \leq 2\text{Reach}(O)$, since $2\text{Reach}(O) \leq \mathcal{D}(O)$ by considering a point on $Sk(O)$, we get:

$$P_O(\mathcal{B}(x, r)) \geq r^d \left(\frac{\text{Reach}(O)}{\mathcal{D}(O)}\right)^d \frac{\omega_d}{\text{Leb}_d(O)},$$

which is also true for r in $[2\text{Reach}(O), \mathcal{D}(O)]$, whereas for $r \geq \mathcal{D}(O)$ we have $P_O(\mathcal{B}(x, r)) = 1$. The choice of a in the lemma is thus relevant.

The case of (a, b) -standard measures

Proof of Proposition 2.28

Before proving Proposition 2.28, we need to prove some lemmas.

Let P be a Borel probability measure supported on a connected compact subset \mathcal{X} of \mathbb{R}^d . We assume this measure to be (a, b) -standard for some positive numbers a and b . In this part, we derive rates of convergence in probability and in expectation for the quantity $\|d_{\hat{P}_N,h} - d_{P,h}\|_{\infty,\mathcal{X}}$. Thanks to these results, we can derive upper bounds and rates of convergence in expectation for $W_1(\mathcal{L}_{N,n,h}(P, P), \mathcal{L}_{N,n,h}^*(\hat{P}_N, \hat{P}_N))$. We finally propose a choice for the parameter N depending on n for which the weak convergences $\mathcal{L}_{N,n,h}(P, P) \rightsquigarrow \|\mathbb{G}_{P,h} - \mathbb{G}'_{P,h}\|_1$ and $\mathcal{L}_{N,n,h}^*(\hat{P}_N, \hat{P}_N) \rightsquigarrow \|\mathbb{G}_{P,h} - \mathbb{G}'_{P,h}\|_1$ occur.

Upper bounds for $\mathbb{P}(\sqrt{n}\|d_{\hat{P}_{N,h}} - d_{P,h}\|_{\infty,\mathcal{X}} \geq \lambda)$ We use the bounds given in Theorem 1 of [CMM16], with the bound for the modulus of continuity given by Lemma 3 in [CMM16]: $\omega(h) = \left(\frac{h}{a}\right)^{\frac{1}{b}}$. We directly get the following lemma:

Lemma 2.43 (Upper bound for $|d_{\hat{P}_{N,h}}(x) - d_{P,h}(x)|$). *Let x be a fixed point in \mathcal{X} and λ a positive number. We have,*

$$\frac{1}{2}\mathbb{P}(|d_{\hat{P}_{N,h}}(x) - d_{P,h}(x)| \geq \lambda) \leq \exp\left(-2a^{\frac{2}{b}}Nh^{\frac{2b-2}{b}}\lambda^2\right) + \exp\left(-\frac{a}{2^{b-1}}N^{\frac{b+1}{2}}h^b\lambda^b\right) + \exp\left(-a^{\frac{1}{b}}N^{\frac{b+1}{2b}}h\lambda\right).$$

In order to derive an upper bound for $\|d_{\hat{P}_{N,h}} - d_{P,h}\|_{\infty,\mathcal{X}}$, like in [CMM16], we use the fact that the function distance-to-measure is 1-Lipschitz and that \mathcal{X} is compact, which means that we can compute a bound by upper-bounding the difference $|d_{\hat{P}_{N,h}}(x) - d_{P,h}(x)|$ over a finite number of points x of \mathcal{X} . Thanks to the following lemma, the minimal number of points needed for this purpose is not bigger than $\frac{(4\mathcal{D}_P\sqrt{d}+\lambda)^d}{\lambda^d}$:

Lemma 2.44. *Let P is a measure supported on \mathcal{X} a compact subset of \mathbb{R}^d , and for $\lambda > 0$ denote $N(P, \lambda) = \inf\{N \in \mathbb{N}, \exists x_1, x_2 \dots x_N \in \mathcal{X}, \bigcup_{i \in [1, N]} B(x_i, \lambda) \supset \mathcal{X}\}$. Then, we have:*

$$N(P, \lambda) \leq \frac{(\mathcal{D}_P\sqrt{d} + \lambda)^d}{\lambda^d}.$$

Proof of Lemma 2.44. The idea is to put a grid on the hypercube containing \mathcal{X} with edges of length \mathcal{D}_P . The grid is a union of small hypercubes with edges of length equal to $\frac{\lambda}{\sqrt{d}}$, so that the number of such small hypercubes into which the big one is split is not superior to $\left(\frac{\mathcal{D}_P\sqrt{d}}{\lambda} + 1\right)^d$.

Then, we decide that each time the intersection between \mathcal{X} and some small hypercube is non-empty, we keep one of the elements of the intersection. We denote x_i the element associated to the i -th hypercube. Finally, each point x in \mathcal{X} belongs to a small hypercube, and its distance to the corresponding x_i is smaller than $\sqrt{\sum_{k=1}^d \frac{\lambda^2}{d}} = \lambda$. \square

We thus derive upper bounds for $\sqrt{n}\|d_{\hat{P}_{N,h}} - d_{P,h}\|_{\infty,\mathcal{X}}$:

Proposition 2.45 (Upper bound for $\sqrt{n}\|d_{\hat{P}_{N,h}} - d_{P,h}\|_{\infty,\mathcal{X}}$). *We have,*

$$\frac{\lambda^d}{2\left(4\mathcal{D}_P\sqrt{d} + \lambda\right)^d} \mathbb{P}(\sqrt{n}\|d_{\hat{P}_{N,h}} - d_{P,h}\|_{\infty,\mathcal{X}} \geq \lambda) \leq \exp\left(-\frac{a^{\frac{2}{b}}Nh^{\frac{2b-2}{b}}}{2}\lambda^2\right) + \exp\left(-\frac{a}{2^{2b-1}}\frac{N^{\frac{b+1}{2}}h^b}{n^{\frac{b}{2}}}\lambda^b\right) + \exp\left(-\frac{a^{\frac{1}{b}}N^{\frac{b+1}{2b}}h}{2n^{\frac{1}{2}}}\lambda\right).$$

Proof of Proposition 2.45. Since the function distance-to-measure is 1-Lipschitz, we get that:

$$\|d_{\hat{P}_{N,h}} - d_{P,h}\|_{\infty,\mathcal{X}} \leq \frac{\lambda}{2} + \sup_i \{|d_{\hat{P}_{N,h}}(x_i) - d_{P,h}(x_i)|\},$$

for the family $(x_i)_i$ associated to a grid which sides are of length equal to $\frac{\lambda}{4\sqrt{d}}$. We can thus bound the probability $\mathbb{P}(\|d_{\hat{P}_{N,h}} - d_{P,h}\|_{\infty,\mathcal{X}} \geq \lambda)$ by

$$\sum_{i=1}^{N(P, \frac{\lambda}{4})} \mathbb{P}\left(|d_{\hat{P}_{N,h}}(x_i) - d_{P,h}(x_i)| \geq \frac{\lambda}{2}\right),$$

with $N(P, \frac{\lambda}{4}) \leq \frac{(4D_P\sqrt{d}+\lambda)^d}{\lambda^d}$ thanks to Lemma 2.44. \square

Upper bounds for the expectation $\mathbb{E}[\|d_{\hat{P}_{N,h}} - d_{P,h}\|_{\infty, \mathcal{X}}]$ In order to get upper bounds for $\mathbb{E}[\|d_{\hat{P}_{N,h}} - d_{P,h}\|_{\infty, \mathcal{X}}]$, we use the same trick as used in [CMM16], which is:

Lemma 2.46. *Let X a random variable such that:*

$$\mathbb{P}(X \geq \lambda) \leq 1 \wedge D\lambda^{-q} \exp(-c\lambda^s)$$

for some integers q and s and some $D > 0$.

We have:

$$\mathbb{E}[X] \leq \left(\frac{\ln c}{c}\right)^{\frac{1}{s}} \left(\frac{q}{s}\right)^{\frac{1}{s}} \left[1 + D \left(\frac{q}{s}\right)^{\frac{-q-s}{s}} \frac{(\ln c)^{\frac{-q-s}{s}}}{s}\right].$$

More particularly, if $c \geq \exp D^{\frac{s}{q+s}} \frac{s}{q}$, then:

$$\mathbb{E}[X] \leq 2 \left(\frac{\ln c}{c}\right)^{\frac{1}{s}} \left(\frac{q}{s}\right)^{\frac{1}{s}}.$$

Proof of Lemma 2.46. For any $\lambda_0 > 0$, that we can choose as $\lambda_0 = \frac{[\ln K]^{\frac{1}{s}}}{c^{\frac{1}{s}}}$, we get that:

$$\begin{aligned} \mathbb{E}[X] &\leq \lambda_0 + \int_{\lambda_0}^{\infty} D\lambda^{-q} \exp(-c\lambda^s) d\lambda \\ &\leq \lambda_0 + D \frac{\lambda_0^{-q-s+1}}{cs} \exp -c\lambda_0^s \\ &= \frac{[\ln K]^{\frac{1}{s}}}{c^{\frac{1}{s}}} + D \frac{[\ln K]^{\frac{-q-s+1}{s}}}{scc^{\frac{-q-s+1}{s}}} \frac{1}{K} \\ &= \frac{[\ln K]^{\frac{1}{s}}}{c^{\frac{1}{s}}} + \frac{[\ln K]^{\frac{1}{s}}}{c^{\frac{1}{s}}} D \frac{[\ln K]^{\frac{-q-s}{s}}}{sc^{\frac{-q}{s}}} \frac{1}{K} \\ &= \frac{[\ln K]^{\frac{1}{s}}}{c^{\frac{1}{s}}} \left[1 + D \frac{[\ln K]^{\frac{-q-s}{s}}}{sKc^{\frac{-q}{s}}}\right] \end{aligned}$$

Finally, if we choose $K = c^{\frac{q}{s}}$, we get:

$$\mathbb{E}[X] \leq \left(\frac{q}{s}\right)^{\frac{1}{s}} \left[\frac{\ln c}{c}\right]^{\frac{1}{s}} \left[1 + D \left[\frac{q}{s}\right]^{\frac{-q-s}{s}} \frac{(\ln c)^{\frac{-q-s}{s}}}{s}\right].$$

\square

From this lemma, we can derive the following lemma.

Lemma 2.47. *We have,*

$$\begin{aligned} \mathbb{E}[\sqrt{n}\|d_{\hat{P}_{N,h}} - d_{P,h}\|_{\infty, \mathcal{X}}] &\leq \\ &\square'_1 \frac{n^{\frac{1}{2}}}{N^{\frac{1}{2}} h^{\frac{b-1}{b}}} \left(\log \left(\frac{Nh^{\frac{2b-2}{b}}}{n}\right)\right)^{\frac{1}{2}} + \\ &\square'_2 \frac{n^{\frac{1}{2}}}{N^{\frac{b+1}{2b}} h} \left(\log \left(\frac{N^{\frac{b+1}{2}} h^b}{n^{\frac{b}{2}}}\right)\right)^{\frac{1}{b}} + \\ &\square'_3 \frac{n^{\frac{1}{2}}}{N^{\frac{b+1}{2b}} h} \log \left(\frac{N^{\frac{b+1}{2b}} h}{n^{\frac{1}{2}}}\right). \end{aligned}$$

for some constants \square depending on a and b .

Proof of part 2 of Proposition 2.28 For all $\lambda > 0$, for any (a, b) -standard measure P supported on a connected compact subset of \mathbb{R}^d , we can use Lemma 2.37 and Lemma 2.47 together with the rates of convergence of the L_1 -Wasserstein distance between empirical and true distribution in [BL14] to get the following result.

If $h \geq \frac{1}{2}$, then for n big enough we have, for some constants \square depending on a and b :

$$\begin{aligned} & \mathbb{E} \left[W_1 \left(\mathcal{L}_{N,n,h}(P, P), \mathcal{L}_{N,n,h}^*(\hat{P}_N, \hat{P}_N) \right) \right] \leq \\ & \square'_1 \frac{n^{\frac{1}{2}}}{(N)^{\frac{1}{2}} h^{\frac{b-1}{b}}} \left(\log \left(\frac{N h^{\frac{2b-2}{b}}}{n} \right) \right)^{\frac{1}{2}} \\ & + \square'_2 \frac{n^{\frac{1}{2}}}{(N)^{\frac{b+1}{2b}} h} \left(\log \left(\frac{N^{\frac{b+1}{2}} h^b}{n^{\frac{b}{2}}} \right) \right)^{\frac{1}{b}} \\ & + \square'_3 \frac{n^{\frac{1}{2}}}{(N)^{\frac{b+1}{2b}} h} \log \left(\frac{N^{\frac{b+1}{2b}} h}{n^{\frac{1}{2}}} \right) \\ & + \square'_4 \frac{n^{\frac{1}{2}}}{N^{\frac{1}{2}}}. \end{aligned}$$

Proof of part 1 of Proposition 2.28 In order to get these two results, we use Lemma 2.24. The convergence to zero of $\sqrt{n} \mathbb{E}[\|d_{P,h} - d_{\hat{P}_{N-n,h}}\|_{\infty, \mathcal{X}}]$ is a direct consequence of Lemma 2.47. We can derive a bound of its rate of convergence in $n^{\frac{1}{2} - \frac{\rho}{2}}$, up to a logarithm term. The a.e. convergence of $\sqrt{n} W_1(d_{P,h}(P), d_{P,h}(\hat{P}_N))$ to zero is derived as in the proof of Proposition 2.26, with the assumption $\rho > 1$. Finally, the a.e. convergence of $\sqrt{n} \|d_{P,h} - d_{\hat{P}_{N,h}}\|_{\infty, \mathcal{X}}$ to zero is a consequence of Proposition 2.45 and of the Borel–Cantelli lemma. It occurs under the assumption $\rho > 1$.

Proof of Theorem 2.29

Lemma 2.48. Let α, κ be two positive numbers and \mathcal{L} and \mathcal{L}^* two laws of real random variables. We denote q_α (respectively q_α^*) the α -quantile of the law \mathcal{L} (respectively \mathcal{L}^*). If $W_1(\mathcal{L}, \mathcal{L}^*) < \kappa$ then:

$$q_\alpha^* \leq 2 \frac{\kappa}{\alpha} + q_{\frac{\alpha}{2}}.$$

Proof. With a drawing, since the L_1 -norm between $F_{\mathcal{L}}$ and $F_{\mathcal{L}^*}$ is smaller than κ , we have:

$$F_{\mathcal{L}^*} \left(q_{\frac{\alpha}{2}} + 2 \frac{\kappa}{\alpha} \right) > 1 - \alpha.$$

□

In this part we assume that h is fixed in $[0, 1]$ and $N = cn^\rho$ for some $\rho > 1$ and $c > 0$. Recall that our aim is to bound above the type II error, that is:

$$\mathbb{P}_{(P,Q)} \left(\sqrt{n} W_1 \left(d_{\hat{P}_N,h}(\hat{P}_n), d_{\hat{Q}_N,h}(\hat{Q}_n) \right) < \hat{q}_\alpha \right).$$

For some $\kappa = n^\gamma$ with γ in $[0, \frac{1}{2})$ to be chosen later, we first bound above the quantile \hat{q}_α with high probability.

As noticed in the proof of Lemma 2.24, the law of $\sqrt{n}W_1\left(d_{P,h}(\hat{P}_n), d_{P,h}(\hat{P}'_n)\right)$ converges to $\mathcal{L}(\|\mathbb{G}_{P,h} - \mathbb{G}'_{P,h}\|_1)$, there is also the convergence of the first moments. So, for n big enough, we have:

$$W_1\left(\mathcal{L}\left(\sqrt{n}W_1\left(d_{P,h}(\hat{P}_n), d_{P,h}(\hat{P}'_n)\right)\right), \mathcal{L}(\|\mathbb{G}_{P,h} - \mathbb{G}'_{P,h}\|_1)\right) \leq 1.$$

Then, under the assumption

$$W_1\left(\mathcal{L}\left(\sqrt{n}W_1\left(d_{P,h}(\hat{P}_n), d_{P,h}(\hat{P}'_n)\right)\right), \mathcal{L}_{N,n,h}^*(\hat{P}_N, \hat{P}_N)\right) \leq \kappa,$$

we have

$$W_1\left(\mathcal{L}(\|\mathbb{G}_{P,h} - \mathbb{G}'_{P,h}\|_1), \mathcal{L}_{N,n,h}^*(\hat{P}_N, \hat{P}_N)\right) \leq \kappa + 1.$$

We can do the same thing for Q . Thus we get that for n big enough and under the previous assumptions:

$$W_1\left(\frac{1}{2}\mathcal{L}(\|\mathbb{G}_{P,h} - \mathbb{G}'_{P,h}\|_1) + \frac{1}{2}\mathcal{L}(\|\mathbb{G}_{Q,h} - \mathbb{G}'_{Q,h}\|_1), \frac{1}{2}\mathcal{L}_{N,n,h}^*(\hat{P}_N, \hat{P}_N) + \frac{1}{2}\mathcal{L}_{N,n,h}^*(\hat{Q}_N, \hat{Q}_N)\right)$$

is bounded above with $\kappa + 1$. And thanks to Lemma 2.48,

$$\hat{q}_\alpha \leq \tilde{q}_{\frac{\alpha}{2}} + 2\frac{\kappa + 1}{\alpha},$$

with \tilde{q}_α the α -quantile of the law $\frac{1}{2}\mathcal{L}(\|\mathbb{G}_{P,h} - \mathbb{G}'_{P,h}\|_1) + \frac{1}{2}\mathcal{L}(\|\mathbb{G}_{Q,h} - \mathbb{G}'_{Q,h}\|_1)$.

We need to notice that with similar arguments as for Lemma 2.37, we have:

$$\begin{aligned} & W_1\left(\mathcal{L}\left(\sqrt{n}W_1\left(d_{P,h}(\hat{P}_n), d_{P,h}(\hat{P}'_n)\right)\right), \mathcal{L}_{N,n,h}^*(\hat{P}_N, \hat{P}_N)\right) \leq \\ & 2\sqrt{n}\mathcal{D}_P\|F_{d_{P,h}(P)} - F_{d_{P,h}(\hat{P}_N)}\|_{\infty,(0,\mathcal{D}_P)} + 2\frac{\sqrt{n}}{h}W_1(P, \hat{P}_N). \end{aligned}$$

Now notice that

$$\begin{aligned} & \sqrt{n}W_1\left(d_{\hat{P}_N,h}(\hat{P}_n), d_{\hat{Q}_N,h}(\hat{Q}_n)\right) \geq \sqrt{n}W_1(d_{P,h}(P), d_{Q,h}(Q)) \\ & - \sqrt{n}W_1\left(d_{\hat{P}_N,h}(\hat{P}_n), d_{P,h}(P)\right) - \sqrt{n}W_1\left(d_{\hat{Q}_N,h}(\hat{Q}_n), d_{Q,h}(Q)\right), \end{aligned}$$

but as well, thanks to Lemma 2.41, the definition of the L_1 -Wasserstein distance as the L_1 -norm between the cumulative distribution functions and to Proposition 2.2:

$$\begin{aligned} & \sqrt{n}W_1\left(d_{\hat{P}_N,h}(\hat{P}_n), d_{P,h}(P)\right) \leq \\ & \frac{\sqrt{n}}{h}W_1(P, \hat{P}_N) + \sqrt{n}\mathcal{D}_{P,h}\|F_{d_{P,h}(\hat{P}_n)} - F_{d_{P,h}(P)}\|_{\infty,(0,\mathcal{D}_P)}, \end{aligned}$$

with $\mathcal{D}_{P,h}$ the diameter of the support of the measure $d_{P,h}(P)$. So, we can finally upper bound $\mathbb{P}_{(P,Q)}\left(\sqrt{n}W_1\left(d_{\hat{P}_N,h}(\hat{P}_n), d_{\hat{Q}_N,h}(\hat{Q}_n)\right) < \hat{q}_\alpha\right)$ by

$$\begin{aligned} & \mathbb{P}\left(\sqrt{n}\mathcal{D}_P\|F_{d_{P,h}(P)} - F_{d_{P,h}(\hat{P}_N)}\|_{\infty,(0,\mathcal{D}_P)} \geq \frac{\kappa}{4}\right) + \\ & \mathbb{P}\left(\sqrt{n}\mathcal{D}_Q\|F_{d_{Q,h}(Q)} - F_{d_{Q,h}(\hat{Q}_N)}\|_{\infty,(0,\mathcal{D}_Q)} \geq \frac{\kappa}{4}\right) + \\ & 2\mathbb{P}\left(\frac{\sqrt{n}}{h}W_1(P, \hat{P}_N) \geq \frac{\kappa}{4}\right) + 2\mathbb{P}\left(\frac{\sqrt{n}}{h}W_1(Q, \hat{Q}_N) \geq \frac{\kappa}{4}\right) + \\ & \mathbb{P}\left(\|F_{d_{P,h}(\hat{P}_n)} - F_{d_{P,h}(P)}\|_{\infty,(0,\mathcal{D}_P)} \geq \frac{W_1(d_{P,h}(P), d_{Q,h}(Q))}{2\mathcal{D}_{P,h}} - \frac{\tilde{q}_{\frac{\alpha}{2}}}{2\mathcal{D}_{P,h}\sqrt{n}} - \frac{(4+\alpha)\kappa+4}{4\mathcal{D}_{P,h}\alpha\sqrt{n}}\right) + \\ & \mathbb{P}\left(\|F_{d_{Q,h}(\hat{Q}_n)} - F_{d_{Q,h}(Q)}\|_{\infty,(0,\mathcal{D}_Q)} \geq \frac{W_1(d_{P,h}(P), d_{Q,h}(Q))}{2\mathcal{D}_{Q,h}} - \frac{\tilde{q}_{\frac{\alpha}{2}}}{2\mathcal{D}_{Q,h}\sqrt{n}} - \frac{(4+\alpha)\kappa+4}{4\mathcal{D}_{Q,h}\alpha\sqrt{n}}\right). \end{aligned}$$

For all positive ϵ , for n big enough, note that the sum of the last two terms can be bounded thanks to the DKW-Massart inequality [Mas90], by

$$4 \exp\left(-\frac{W_1^2(d_{P,h}(P), d_{Q,h}(Q))}{(2+\epsilon)\max\{\mathcal{D}_P^2, \mathcal{D}_Q^2\}}n\right).$$

Note also that thanks to the DKW-Massart inequality, the first term can be bounded above by

$$2 \exp\left(-\frac{1}{8\mathcal{D}_P^2}cn^{\rho-1+2\gamma}\right).$$

The second term is similar. Thanks to Theorem 2 in [FG15], the third term is bounded above by

$$c_1 \exp\left(-c_2 h^d n^{\rho+d\gamma-\frac{d}{2}}\right),$$

for some fixed constants c_1 and c_2 . The remaining terms are similar.

Since $\rho > 1$, we can choose a positive γ satisfying: $\gamma < \frac{1}{2}$, $\rho + d\gamma - \frac{d}{2} > 1$ and $\rho - 1 + 2\gamma > 1$. So the two last expressions are negligible in comparison to the first one.

So, for n big enough, $\mathbb{P}_{(P,Q)}\left(\sqrt{n}W_1\left(d_{\hat{P}_N,h}(\hat{P}_n), d_{\hat{Q}_N,h}(\hat{Q}_n)\right) < \hat{q}_\alpha\right)$ is bounded above by

$$4 \exp\left(-\frac{W_1^2(d_{P,h}(P), d_{Q,h}(Q))}{3\max\{\mathcal{D}_{P,h}^2, \mathcal{D}_{Q,h}^2\}}n\right).$$

Appendix

2.A A characterization of metric-measure spaces up to an isomorphism – On the proof of a theorem from Gromov

Le Chapitre 2.A présente une preuve détaillée du théorème de mm-Reconstruction de Gromov [GLP99, Theorem 3 1/2.5.]. Dans ces travaux, Gromov se restreint à des espaces Polonais munis d'une mesure de probabilité. Il propose une caractérisation de l'ensemble des espaces métriques mesurés ainsi obtenus, à isomorphisme près. Le théorème énonce le fait suivant : la connaissance, pour tout r , de la distribution de la matrice des distances associée à un r -échantillon tiré selon une mesure suffit à reconstruire l'espace métrique mesuré associé, à isomorphisme près. Chapter 2.A presents a detailed proof of the mm-Reconstruction theorem from Gromov [GLP99, Theorem 3 1/2.5.]. In this work, Gromov studies Polish spaces equipped with a probability distribution. He offers a characterization of the set of metric measure spaces obtained in this way, up to an isomorphism. The theorem states the following fact : the knowledge, for all r , of the distribution of the matrix of distances associated to a r -sample from a measure suffices to reconstruct the associated metric measure space, up to an isomorphism.

2.A.1 The Theorem of identifiability of measures up to an isomorphism

In this chapter, (\mathcal{X}, δ) is a Polish space that we equip with a Borel probability measure $P \in \mathcal{M}(\mathcal{X})$. For all $r \in \mathbb{N}^*$ we define the distribution $P_r = \psi_r(P)$ which is the push-forward of P by the application $\psi_r : \mathcal{X}^r \rightarrow \text{Mat}_r(\mathbb{R}_+)$, $(x_1, x_2, \dots, x_r) \mapsto (\delta(x_i, x_j))_{i,j \in [1,r]}$. Here, $\text{Mat}_r(\mathbb{R}_+)$ denotes the set of $r \times r$ matrices with elements in \mathbb{R}_+ .

The Gromov theorem states as follows.

Theorem 2.49 (mm-Reconstruction theorem [GLP99]). *Let (\mathcal{X}, δ, P) and (\mathcal{Y}, γ, Q) be two Polish mm-spaces such that P is supported on \mathcal{X} and Q supported on \mathcal{Y} . Then the following property is satisfied. If the measures P_r and Q_r are equal for all $r \in \mathbb{N}^*$, then (\mathcal{X}, δ, P) and (\mathcal{Y}, γ, Q) are isomorphic.*

The proof of this theorem relies on the well-know theorem of isometries extension.

Lemma 2.50 (Isometries extension theorem). *Let (\mathcal{X}, δ) and (\mathcal{Y}, γ) be two metric spaces, with (\mathcal{Y}, γ) complete. Let A be a dense subset of \mathcal{X} and $f : A \rightarrow \mathcal{Y}$ an isometry. Then, there exists a unique isometry $g : \mathcal{X} \rightarrow \mathcal{Y}$ which coincides with f on A .*

Proof. For all sequence $(x_n)_{n \in \mathbb{N}}$ in A converging to $x \in \mathcal{X}$, $(f(x_n))_{n \in \mathbb{N}}$ is Cauchy in the complete space \mathcal{Y} since f is an isometry, thus $(f(x_n))_{n \in \mathbb{N}}$ converges to some limit $g(x) \in \mathcal{Y}$. Note that this limit does not depend on the sequence $(x_n)_{n \in \mathbb{N}}$. The map g built in this way is an isometry. \square

Let (\mathcal{X}, δ, P) and (\mathcal{Y}, γ, Q) be two Polish mm-spaces such that the measures P_r and Q_r coincide for all $r \in \mathbb{N}^*$. Proving Theorem 2.49 boils down to build an isomorphism between both mm-spaces. We proceed as follows. First we define an isomorphism $\psi : x_i^* \mapsto y_i^*$ on a dense countable subset of \mathcal{X} , $\{x_i^* \mid i \in \mathbb{N}\}$. Then, according to Lemma 2.50, it will be possible to extend ψ to an isomorphism $\tilde{\psi}$ defined on the whole space \mathcal{X} . We also need to check that $\tilde{\psi}$ is one-to-one and onto. Finally, we prove that the map $\tilde{\psi}$ preserves the measures and is thus an isomorphism between (\mathcal{X}, δ, P) and (\mathcal{Y}, γ, Q) .

Let $(x_i^*)_{i \in \mathbb{N}}$ be a dense sequence in \mathcal{X} . Such a sequence exists since (\mathcal{X}, δ) is a Polish space, thus separable. The set R_1 defined by

$$R_1 = \bigcup_{i \in \mathbb{N}} \{\rho \in \mathbb{R}^+ \mid P(\partial \mathcal{B}(x_i^*, \rho)) > 0\}$$

is a countable subset of \mathbb{R}^+ , as a countable union of countable sets. Then we can set $R = (\rho_n)_{n \in \mathbb{N}}$, a sequence of elements in $\mathbb{R}^{+,*} \setminus R_1$ which is dense in \mathbb{R}^+ . Since the set $\bar{R}^\infty = R \cup \{+\infty\}$ is countable, the set of maps $\bigcup_{n \in \mathbb{N}} \{\phi : \llbracket 1, n \rrbracket \rightarrow \bar{R}^\infty\}$ is also countable. Thus, we can define a sequence $(\phi_n)_{n \in \mathbb{N}^*}$ which contains all of the elements of $\bigcup_{n \in \mathbb{N}} \{\phi : \llbracket 1, n \rrbracket \rightarrow \bar{R}^\infty\}$. For convenience, for all $n \in \mathbb{N}^*$, we extend ϕ_n to \mathbb{N} with $\phi_n(i) = +\infty$ whenever i is not in the definition domain of ϕ_n .

We define the map $\psi_n : \mathcal{X}^n \rightarrow \mathbb{R}_+^{\frac{3n^2+n}{2}}$ for all $x^n = (x_1, \dots, x_n) \in \mathcal{X}^n$ by:

$$\psi_n(x^n) = \left((P(\mathcal{B}(x_i, \phi_j(i))))_{i,j \in \llbracket 1, n \rrbracket}, \left(P \left(\bigcap_{i \in \llbracket 1, l \rrbracket} \mathcal{B}(x_i, \phi_j(i)) \right) \right)_{j,l \in \llbracket 1, n \rrbracket}, (\delta(x_i, x_j))_{i,j \in \llbracket 1, n \rrbracket, i < j} \right). \quad (2.7)$$

According to Lemma 2.51, for all $n \in \mathbb{N}$, $\psi_n(P^{\otimes n})$ and $\psi_n(Q^{\otimes n})$ coincide.

Lemma 2.51. *If $P_r = Q_r$ for all $r \in \mathbb{N}$, then for all $n \in \mathbb{N}^*$ $\psi_n(P^{\otimes n}) = \psi_n(Q^{\otimes n})$, with ψ_n defined in (2.7).*

Set $(y^n)_{n \in \mathbb{N}}$, a sequence of elements $y^n \in \mathcal{Y}^{\mathbb{N}}$ obtained as in Lemma 2.52.

Lemma 2.52. *We can build a sequence $(y^n)_{n \in \mathbb{N}}$ of elements $y^n = (y_{1,n}, y_{2,n}, \dots, y_{n,n}, a, \dots, a, \dots) \in \mathcal{Y}^{\mathbb{N}}$ with a fixed point in \mathcal{Y} , such that $\forall n \in \mathbb{N}$:*

- $\forall i, j \in \llbracket 1, n \rrbracket, |P(\mathcal{B}(x_i^*, \phi_j(i))) - Q(\mathcal{B}(y_{i,n}, \phi_j(i)))| \leq \frac{1}{n}$,
- $\forall j, l \in \llbracket 1, n \rrbracket, |P(\bigcap_{i \in \llbracket 1, l \rrbracket} \mathcal{B}(x_i^*, \phi_j(i))) - Q(\bigcap_{i \in \llbracket 1, l \rrbracket} \mathcal{B}(y_{i,n}, \phi_j(i)))| \leq \frac{1}{n}$,
- $\forall i, j \in \llbracket 1, n \rrbracket, |\delta(x_i^*, x_j^*) - \gamma(y_{i,n}, y_{j,n})| \leq \frac{1}{n}$.

With a diagonal process, see Lemma 2.53, we can extract a subsequence of $(y^n)_{n \in \mathbb{N}}$ such that $y_{i,n}$ converges to a point in \mathcal{Y} , which we denote by y_i^* .

Lemma 2.53. *We can build a subsequence of $(y^n)_{n \in \mathbb{N}}$ such that for all $i \in \mathbb{N}^*$, the i -th coordinate converges to a point y_i^* in \mathcal{Y} .*

According to Lemma 2.54, the points y_i^* satisfy:

$$Q \left(\bigcap_{i \in \llbracket 1, n \rrbracket} \mathcal{B}(y_i^*, \phi_j(i)) \right) = P \left(\bigcap_{i \in \llbracket 1, n \rrbracket} \mathcal{B}(x_i^*, \phi_j(i)) \right), \quad (2.8)$$

and $\delta(x_i^*, x_j^*) = \gamma(y_i^*, y_j^*)$.

Lemma 2.54. *The points $(y_i^*)_{i \in \mathbb{N}^*}$ satisfy:*

$$(i) \quad \forall j, n \in \mathbb{N}^*, Q \left(\bigcap_{i \in \llbracket 1, n \rrbracket} \mathcal{B}(y_i^*, \phi_j(i)) \right) = P \left(\bigcap_{i \in \llbracket 1, n \rrbracket} \mathcal{B}(x_i^*, \phi_j(i)) \right)$$

$$(ii) \quad \forall i, j \in \mathbb{N}^*, \delta(x_i^*, x_j^*) = \gamma(y_i^*, y_j^*).$$

According to Lemma 2.54, the map ψ defined on $(x_i^*)_{i \in \mathbb{N}}$ by $\psi(x_i^*) = y_i^*$ is an isometry. Since (\mathcal{Y}, γ) is complete, Lemma 2.50 entails that ψ can be extended to an isometry $\tilde{\psi}$ on \mathcal{X} which coincides with ψ on $(x_i^*)_{i \in \mathbb{N}}$.

Moreover, Lemma 2.55 yields that the sets $\psi(\mathcal{X})$ and \mathcal{Y} are equal. Thus, $\tilde{\psi}$ is an isometry from \mathcal{X} to \mathcal{Y} . In addition, we get that $(y_i^*)_{i \in \mathbb{N}}$ is dense in \mathcal{Y} .

Lemma 2.55 (Equality of $\psi(\mathcal{X})$ and \mathcal{Y}). *The sets $\psi(\mathcal{X})$ and \mathcal{Y} are equal.*

Note that the collection $\{\bigcap_{i \in \llbracket 1, n \rrbracket} (\mathcal{B}(y_i^*, \phi_j(i))) \mid j, n \in \mathbb{N}^*\}$ is a π -system generating the σ -algebra of Borel sets of \mathcal{Y} . Indeed, we can write any open set as a countable union of balls of radius in $\overline{\mathbb{R}}^\infty$ centred on the y_i^* s. Moreover, the measures Q and $\psi(P)$ coincide on this π -system. Indeed, according to Lemma 2.54, for all $j, n \in \mathbb{N}^*$

$$\begin{aligned} Q \left(\bigcap_{i \in \llbracket 1, n \rrbracket} \mathcal{B}(y_i^*, \phi_j(i)) \right) &= P \left(\bigcap_{i \in \llbracket 1, n \rrbracket} \mathcal{B}(x_i^*, \phi_j(i)) \right) \\ &= P \left(\left\{ y \in \mathcal{X} \mid \psi(y) \in \bigcap_{i \in \llbracket 1, n \rrbracket} \mathcal{B}(y_i^*, \phi_j(i)) \right\} \right). \end{aligned}$$

Thus, according to the monotone class lemma, the measures Q and $\psi(P)$ coincide. This concludes the proof of Theorem 2.49.

2.A.2 Proofs

Proof of Lemma 2.51

The image of the map ψ_n is contained in \mathbb{R}_+^d with $d = \frac{5n^2 - n}{2}$. According to the monotone class lemma, proving equality of the probability measures $\psi_n(P^{\otimes n})$ and $\psi_n(Q^{\otimes n})$ boils down to prove that they coincide on the π -system $\{]a_1, b_1[\times \dots \times]a_d, b_d[\mid (a_i, b_i) \in \mathbb{Q}^2, a_i < b_i\}$.

Set $d' = 2n^2$, we approximate $\mathbb{1}_{]a_i, b_i[}$ for all $i \in \llbracket 1, d' \rrbracket$ with a bounded continuous function $\alpha_{l,i}$ which is equal to 1 on $]a_i - \frac{1}{l}, b_i + \frac{1}{l}[$ and to 0 on $]a_i, b_i[^c$ for l large enough, so that:

$$\lim_{l \rightarrow \infty} \int_{[0,1]^{d'} \times \mathbb{R}_+^{d-d'}} \prod_{i=1}^{d'} \alpha_{l,i}(t_i) \prod_{i=d'+1}^d \mathbb{1}_{]a_i, b_i[}(t_i) d\psi_n(P^{\otimes n})(t_1, \dots, t_d) = \psi_n(P^{\otimes n})(]a_1, b_1[\times \dots \times]a_d, b_d[).$$

It remains to prove equality of the integrals associated to both measures, for all $\alpha_{l,i}$ s. The Stone-Weierstrass theorem entails that for all bounded continuous function α defined on $[0, 1]^{d'}$, the integrals $\int_{[0,1]^{d'} \times \mathbb{R}_+^{d-d'}} \alpha(t_1, \dots, t_{d'}) \prod_{i=d'+1}^d \mathbb{1}_{]a_i, b_i[}(t_i) d\psi_n(P^{\otimes n})(t_1, \dots, t_d)$ coincide for both

measures. Indeed, these integrals coincide for all polynomial function α . For instance, as noted in [GLP99, Theorem 3 1/2.5.], Fubini-Tonelli theorem and the fact that the measures P_r and Q_r coincide for all $r \in \mathbb{N}$ yield

$$\int_{\mathcal{X}} P(\mathcal{B}(x, \rho)) dP(x) = \int_{\mathcal{X}} \int_{\mathcal{X}} \mathbb{1}_{\delta(x,y) \leq \rho} dP(x) dP(y) = \int_{\mathcal{Y}} Q(\mathcal{B}(x, \rho)) dQ(x).$$

This method can be generalised to all products of polynomial functions of measures of balls, of intersections of balls and of functions $x, y \mapsto \delta(x, y)$.

Proof of Lemma 2.52

The following lemma, Lemma 2.56, ensures the continuity of all the maps ψ_n defined by (2.7) for $n \in \mathbb{N}^*$.

Lemma 2.56. *For all $n \in \mathbb{N}^*$, the map $\psi_n : \mathcal{X}^n \rightarrow \mathbb{R}_+^{\frac{5n^2-n}{2}}$ defined in (2.7) is continuous.*

Proof. We equip the space \mathcal{X}^n with the metric δ^n defined for all $x = (x_1, x_2, \dots, x_n)$ and $z = (z_1, z_2, \dots, z_n)$ by $\delta^n(x, z) = \max_{i \in \llbracket 1, n \rrbracket} (\delta(x_i, z_i))$.

Let $(z^k)_{k \in \mathbb{N}} = ((z_1^k, z_2^k, \dots, z_n^k))_{k \in \mathbb{N}}$ be converging sequence in \mathcal{X}^n with limit $x = (x_1, x_2, \dots, x_n)$. Then $|\delta(z_i^k, z_j^k) - \delta(x_i, x_j)| \leq 2\delta^n(x, z^k)$ which converges to zero when k goes to infinity.

Note that $P(\mathcal{B}(z_i^k, \phi_j(i))) \leq P(\mathcal{B}(x_i, \phi_j(i) + \delta^n(x, z^k)))$. Thus, $\limsup_{k \rightarrow \infty} P(\mathcal{B}(z_i^k, \phi_j(i))) \leq P(\overline{\mathcal{B}}(x_i, \phi_j(i)))$. As well, we get that $\liminf_{k \rightarrow \infty} P(\mathcal{B}(z_i^k, \phi_j(i))) \geq P(\mathcal{B}(x_i, \phi_j(i)))$.

Since for all j , ϕ_j is chosen such that $\forall i \in \mathbb{N}$, $P(\partial \mathcal{B}(x_i, \phi_j(i))) = 0$, it holds that

$$P(\mathcal{B}(x_i, \phi_j(i))) = P(\overline{\mathcal{B}}(x_i, \phi_j(i))) = \lim_{k \rightarrow \infty} P(\mathcal{B}(z_i^k, \phi_j(i))).$$

The same result holds for the intersections of balls. Thus, the map ψ_n is continuous. \square

Since $\text{Supp}(P) = \mathcal{X}$, $x^n = (x_1, x_2, \dots, x_n) \in \text{Supp}(P^{\otimes n})$. It yields that for all $\eta > 0$, $P^{\otimes n}(\mathcal{B}(x^n, \eta)) > 0$. Moreover, according to Lemma 2.56, ψ_n is continuous on \mathcal{X}^n . Thus, for all $\epsilon > 0$, there exists $\eta > 0$ such that $\mathcal{B}(x^n, \eta) \subset \{z = (z_1, z_2, \dots, z_n) \in \mathcal{X}^n \mid \|\psi_n(x^n) - \psi_n(z)\|_\infty \leq \epsilon\}$. It holds that $\psi_n(x^n) \in \text{Supp}(\psi_n(P^{\otimes n}))$.

According to Lemma 2.51, the distributions $\psi_n(P^{\otimes n})$ and $\psi_n(Q^{\otimes n})$ coincide. Thus, $\psi_n(x^n) \in \text{Supp}(\psi_n(Q^{\otimes n}))$, and $\forall \epsilon > 0$, $\psi_n(Q^{\otimes n})(\mathcal{B}(\psi_n(x^n), \epsilon)) > 0$.

It yields that for all $n \in \mathbb{N}^*$, there exists some $y^n = (y_{1,n}, y_{2,n}, \dots, y_{n,n}) \in \mathcal{Y}^n = \text{Supp}(Q^{\otimes n})$ which satisfies:

$$\|\psi_n(y^n) - \psi_n(x^n)\|_\infty < \frac{1}{n}.$$

Proof of Lemma 2.53

According to Lemma 2.52, $(y_{i,n})_{n \in \mathbb{N}^*}$ is built such that for all $j \in \llbracket 1, n \rrbracket$, $|P(\mathcal{B}(y_{i,n}, \phi_j(i))) - P(\mathcal{B}(x_i^*, \phi_j(i)))| \leq \frac{1}{n}$. For all $\epsilon > 0$, there exists $N \in \mathbb{N}^*$ such that $0 < \phi_N(i) < \epsilon$. Thus, for all $n \geq N$, $P(\mathcal{B}(y_{i,n}, \epsilon)) \geq P(\mathcal{B}(x_i^*, \phi_N(i))) - \frac{1}{n}$. And for n large enough, $P(\mathcal{B}(y_{i,n}, \epsilon)) \geq \frac{1}{2} P(\mathcal{B}(x_i^*, \phi_N(i)))$ which is positive since $x_i^* \in \text{Supp}(P)$.

Thus, from any subsequence of $(y_{i,n})_{n \in \mathbb{N}^*}$ we extract a converging subsequence, according to Lemma 2.57.

In order to make the sequences $(y_{i,n})_{n \in \mathbb{N}^*}$ for $i \in \mathbb{N}^*$ converge simultaneously, we apply a diagonal process.

Lemma 2.57. *Let $(x_n)_{n \in \mathbb{N}}$ be a sequence in a complete metric space (\mathcal{X}, δ) equipped with a Borel measure P , such that for all $\epsilon > 0$ it exists $c > 0$ and $N > 0$ such that for all $n \geq N$, $P(\mathcal{B}(x_n, \epsilon)) > c$. Then, a converging subsequence can be extracted from $(x_n)_{n \in \mathbb{N}}$.*

Proof. For $\epsilon > 0$, let c and N be defined such that for all $n \geq N$, $P(\mathcal{B}(x_n, \epsilon)) > c$. Set $L(\epsilon) = \sup\{m \in \mathbb{N} \mid \exists n_1, n_2, \dots, n_m \geq N, \forall i, j \in \llbracket 1, m \rrbracket \mathcal{B}(x_{n_i}, \epsilon) \cap \mathcal{B}(x_{n_j}, \epsilon) = \emptyset\}$. It corresponds to the packing number of $\{\mathcal{B}(x_n, \epsilon)\}_{n \geq N}$. Note that $L(\epsilon) \leq \frac{1}{c} < \infty$.

Thus, there is a set of points $(x_l^*)_{l \in \llbracket 1, L(\epsilon) \rrbracket}$ maximizing the previous set. Note that $\mathcal{X} = \bigcup_{l \in \llbracket 1, L(\epsilon) \rrbracket} \mathcal{B}(x_l^*, 2\epsilon)$. Thus, there is a ball of radius 2ϵ containing an infinite subset of points of the sequence $(x_n)_{n \in \mathbb{N}^*}$. We denote any such infinite subset by $\mathcal{S}((x_n)_{n \in \mathbb{N}^*}, \epsilon)$.

To build a subsequence $(z_n)_{n \in \mathbb{N}^*}$, we proceed as follows. First set $z_1 = x_1$ and keep $S_1 = \mathcal{S}((x_n)_{n \in \mathbb{N}^*}, 1) \setminus \{z_1\}$. For $n \in \mathbb{N}^*$, select for z_{n+1} the first element of $S_{n+1} = \mathcal{S}(S_n, \frac{1}{n+1}) \setminus \{z_n\}$. Then $(z_n)_{n \in \mathbb{N}^*}$ is a subsequence of $(x_n)_{n \in \mathbb{N}^*}$ which is Cauchy in a complete metric space, thus converging. \square

Proof of Lemma 2.54

Triangular inequality and Lemma 2.52 yield $\gamma(y_i^*, y_j^*) = \delta(x_i^*, x_j^*)$.

According to Lemma 2.52,

$$\forall j, l \in \llbracket 1, n \rrbracket, \left| P \left(\bigcap_{i \in \llbracket 1, l \rrbracket} \mathcal{B}(x_i^*, \phi_j(i)) \right) - Q \left(\bigcap_{i \in \llbracket 1, l \rrbracket} \mathcal{B}(y_{i,n}, \phi_j(i)) \right) \right| \leq \frac{1}{n}.$$

For all $j, i \in \mathbb{N}$, we can build an increasing sequence $(\phi_{j,k}(i))_{k \in \mathbb{N}}$ converging to $\phi_j(i)$.

For all k, l in \mathbb{N}^* , for m large enough, $\bigcap_{i \in \llbracket 1, l \rrbracket} \mathcal{B}(y_i^*, \phi_j(i)) \supset \bigcap_{i \in \llbracket 1, l \rrbracket} \mathcal{B}(y_{i,m}, \phi_{j,k}(i))$. Thus,

$$\begin{aligned} Q \left(\bigcap_{i \in \llbracket 1, l \rrbracket} \mathcal{B}(y_i^*, \phi_j(i)) \right) &\geq Q \left(\bigcap_{i \in \llbracket 1, l \rrbracket} \mathcal{B}(y_{i,m}, \phi_{j,k}(i)) \right) \\ &\geq P \left(\bigcap_{i \in \llbracket 1, l \rrbracket} \mathcal{B}(x_i^*, \phi_{j,k}(i)) \right) - \frac{1}{m} \end{aligned}$$

By making m and k go to infinity, we get that:

$$P \left(\bigcap_{i \in \llbracket 1, l \rrbracket} \mathcal{B}(x_i^*, \phi_j(i)) \right) \leq Q \left(\bigcap_{i \in \llbracket 1, l \rrbracket} \mathcal{B}(y_i^*, \phi_j(i)) \right).$$

Similarly, we prove that $P \left(\bigcap_{i \in \llbracket 1, l \rrbracket} \overline{\mathcal{B}}(x_i^*, \phi_j(i)) \right) \geq Q \left(\bigcap_{i \in \llbracket 1, l \rrbracket} \mathcal{B}(y_i^*, \phi_j(i)) \right)$. Since the ρ_j s are chosen such that $P \left(\bigcap_{i \in \llbracket 1, l \rrbracket} \overline{\mathcal{B}}(x_i^*, \phi_j(i)) \right) = P \left(\bigcap_{i \in \llbracket 1, l \rrbracket} \mathcal{B}(x_i^*, \phi_j(i)) \right)$, the result follows.

Proof of Lemma 2.55

First note that $\psi(\mathcal{X})$ is a closed subset of \mathcal{Y} by using the completeness of \mathcal{Y} . For the sake of contradiction, assume that $\psi(\mathcal{X}) \neq \mathcal{Y}$ and choose $y \in \mathcal{Y} \setminus \psi(\mathcal{X})$. Then, for some $\epsilon > 0$, $\mathcal{B}(y, \epsilon) \subset \mathcal{Y} \setminus \psi(\mathcal{X})$. Since $y \in \text{Supp}(Q)$, $Q(\mathcal{B}(y, \frac{\epsilon}{2})) > 0$.

Let $\rho \in R$ be such that $\rho < \frac{\epsilon}{2}$, then $Q(\mathcal{Y} \setminus \mathcal{B}(y, \frac{\epsilon}{2})) \geq \lim_{n \rightarrow \infty} Q(\bigcup_{i \leq n} \mathcal{B}(y_i^*, \rho))$. The Poincaré's formula and Lemma 2.54 (with j chosen such that for all $i \in \llbracket 1, n \rrbracket$, $\phi_j(i) = \rho$) yield $P(\bigcup_{i \leq n} \mathcal{B}(x_i^*, \rho)) = Q(\bigcup_{i \leq n} \mathcal{B}(y_i^*, \rho))$. By making n go to infinity, it holds that $1 = P(\mathcal{X}) = Q(\mathcal{Y} \setminus \mathcal{B}(y, \frac{\epsilon}{2}))$, which is a contradiction.

Trimmed criteria for clustering with applications to mixture models

Dans le chapitre précédent il était question de comparer deux distributions ou deux nuages de points. Désormais, nous considérons un unique nuage de points, ou une unique distribution. L'objectif de ce chapitre est d'introduire/étudier des méthodes permettant de séparer une telle distribution ou un tel nuage en k composantes, de façon pertinente. Les méthodes étudiées sont adaptées à des données échantillonnées selon certains modèles de mélange, en présence de données aberrantes. Dans un premier temps, nous étudions une version trimmée à la façon [CAGM97] de la méthode de partitionnement Bregman proposée par [BMDG05]. Ces travaux ont été réalisés *en collaboration avec Aurélie Fischer et Clément Levrard* et ont fait l'objet d'un papier [BFL18] soumis à la conférence *NIPS 2018*. Dans un second temps, nous développons une nouvelle méthode de partitionnement trimmé adaptée aux mélanges Gaussiens hétéroscédastiques. Cette méthode repose sur une version continue de la vraisemblance trimmée. C'est une alternative à la méthode de [GEGMMI08] implémentée par l'algorithme R *tclust*. L'originalité de notre alternative vient de la possible non convexité/connexité des composantes. Cette caractéristique rend notre méthode plus performante pour certains types de distributions. Les travaux associés à cette deuxième partie sont *en cours*.

Contents

3.1	Trimmed clustering with Bregman divergences	79
3.1.1	Context and motivation	80
3.1.2	Contribution	81
3.1.3	A framework for trimmed clustering with Bregman divergences	82
3.1.4	Algorithm	87
3.1.5	Experiments	88
3.2	Heteroscedastic Gaussian mixture clustering with the trimmed log-likelihood	94
3.2.1	The trimmed log-likelihood - a notion related to the distance to measure	95
3.2.2	A new clustering method based on the trimmed log-likelihood and a modified version of the Lloyd's algorithm	99
3.2.3	Experiments	105
3.3	Proofs	109
3.3.1	Proofs for Section 3.1	109
3.3.2	Proofs for Section 3.2	120

In this chapter, we consider a distribution or a sample of points. We aim at splitting the data into k components. The methods we study are well suited for data generated according to certain mixtures of models, in the presence of outliers or clutter noise. On the one hand, we consider a trimmed version (in the sense of [CAGM97]) of the Bregman clustering method introduced in [BMDG05]. On the other hand, we develop a new method of trimmed clustering well-suited to heteroscedastic Gaussian mixture models. This method is based on a continuous version of the trimmed log-likelihood. This is an alternative to the work of [GEGMMI08] with the R algorithm *tclust*. The originality of our point of view resides in the possible non convexity/connectedness of the k components. This property makes our method perform better for some distributions.

3.1 Trimmed clustering with Bregman divergences

In this work we investigate a k -means type method for clustering data possibly corrupted with clutter noise. We propose a trimming approach based on Bregman divergences. The main interest in using Bregman divergences for clustering is two-sided: first it adapts to data sampled according to mixture models from exponential families, and second there is a possibility to apply

Lloyd-type algorithms. In this work we prove that the optimal empirical codebook converges a.e. to the optimal codebook, which is proved to exist. We also prove that the Bregman-loss associated with the empirical optimal codebook approximates the optimal Bregman-loss at a rate $\frac{1}{\sqrt{n}}$, which is the usual rate for k -means.

Moreover, the Bregman-loss associated with some trimming parameter h coincides with a notion of distance-to-measure, where the squared Euclidean distance is replaced by a Bregman divergence in the definition of the DTM. Then, optimal codebooks (resp. empirical codebooks) coincide with minima of this generalization of distance-to-measure (resp. empirical distance-to-measure). This generalization of the DTM will be called Bregman divergence-to-measure.

Finally, we derive a Lloyd-type algorithm with a trimming parameter that can be selected from data according to some heuristic.

3.1.1 Context and motivation

Given a set of data points sampled according to some distribution P on \mathbb{R}^d , possibly corrupted by noise, it is often a question in data analysis to determine clusters together with centres associated to each cluster, for instance to provide a concise description of the data. Given a positive integer k , a set $\mathbf{c} = \{c_1, c_2, \dots, c_k\}$ of k elements in \mathbb{R}^d , or k centres, will be called a *codebook*. We focus on the problem of choosing the codebook that best depicts the data, together with a quantizer, that is a map $q : \mathbb{R}^d \rightarrow \mathbf{c}$ that assigns each element of \mathbb{R}^d to some centre $c_i = q(x)$ in \mathbf{c} . Whenever x is a data point that satisfies $c_i = q(x)$, we say that the *label* of the point x is i . This kind of method consisting in assigning each data point to a single centre is called *hard clustering*. This is in opposition to *soft clustering* where to each data point x we assign, for each center, a number corresponding to the probability for x to belong to the corresponding cluster. Also, we focus on *unsupervised clustering*, in the sense that the data come without labeling, contrary to *supervised clustering* where some of the data points are already labeled.

One of the most famous and widely used hard unsupervised clustering methods is k -means. This method is based on the concept of *variation about \mathbf{c} given P* , $V^P(\mathbf{c})$ defined by

$$V^P(\mathbf{c}) = P \min_{i \in \llbracket 1, k \rrbracket} \|u - c_i\|^2,$$

where we use the notation $\llbracket 1, k \rrbracket$ for the set $\{1, 2, \dots, k\}$, and for any function f , $Pf(u)$ denotes the expectation of the random variable $f(X)$ when X is sampled according to P . The symbol $\|\cdot\|$ denotes the Euclidean norm on \mathbb{R}^d . The *loss* or k -*variation given P* for k -means is given by

$$V_k^P = \inf_{\mathbf{c}} V^P(\mathbf{c}).$$

Any minimizer $\mathbf{c}^* = \{c_1^*, c_2^*, \dots, c_k^*\}$ of the k -variation given P will be called an *optimal codebook*. Then, an *optimal quantizer* associated to the k -means method is given by $q^* : x \mapsto c_{i(x)}^*$ whenever $\|x - c_{i(x)}^*\| = \min_{i \in \llbracket 1, k \rrbracket} \|x - c_i^*\|$. This provides a clustering method, where each data point x is assigned to the cluster i whenever $q^*(x) = c_i^*$. Intuitively, when the form of the data is truly a union of k well-separated clusters, then \mathbf{c}^* should correspond to the set of all barycentres c_i^* of data points within a cluster i , and every point in the cluster i should be assigned to c_i^* though q^* .

When the data are corrupted with outliers, it is more appropriate to decide not to assign all data points to a cluster, that is not to label all data points. For convenience we assign the label 0 to a proportion $1 - h$ of the data points. Such points will be considered as noise and will count for 0 in the remaining loss $V_{k,h}(P)$, also called *trimmed k -variation* which is defined by

$$V_{k,h}(P) = \inf_{\tilde{P} \in \mathcal{P}_{h^+}(P)} V_k^{\tilde{P}},$$

where $\mathcal{P}_{h+}(P) = \bigcup_{h' \in [h,1]} \mathcal{P}_{h'}(P)$ with $\mathcal{P}_h(P)$, the set of probability measures $\tilde{P} = \frac{1}{h}Q$ for Q a sub-measure of P of P -mass h .

Such a point of view has been taken in [CAGM97]; it corresponds to the concept of *trimmed k -means*. It consists of applying a k -means method to a subset of \mathbb{R}^d with P -mass at least h , such that the loss is minimized. In practice, when we observe a set of data points from P , corrupted by clutter noise, a suitable choice of h would be the proportion of signal, that is of data points truly sampled from P . Moreover, we can expect that this set minimizes the loss. Roughly, noisy data points should have 0 as a label.

In [CAGM97], the authors generalise these notions by replacing the squared Euclidean norm $\|\cdot\|^2$ with a convex function of the Euclidean norm $x \mapsto \phi(\|x\|)$, where ϕ is convex on \mathbb{R} , in the definition of the variation about \mathbf{c} given P .

To every codebook \mathbf{c} , it is possible to assign a radius $r_h(\mathbf{c})$, which is the largest radius $r \geq 0$ such that $P(\mathcal{B}(\mathbf{c}, r)) \leq h \leq P(\overline{\mathcal{B}}(\mathbf{c}, r))$ with the notation $\mathcal{B}(\mathbf{c}, r) = \bigcup_{i \in [1,k]} \mathcal{B}(c_i, r)$ and $\overline{\mathcal{B}}(\mathbf{c}, r)$ is its closure, where $\mathcal{B}(c, r) = \{x \in \mathbb{R}^d \mid \|x - c\| < r\}$ denotes the Euclidean ball. The family of distributions $P_{\mathbf{c},h} \in \mathcal{P}_h(P)$ supported on $\overline{\mathcal{B}}(\mathbf{c}, r)$ and such that $hP_{\mathbf{c},h}$ coincides with P on $\mathcal{B}(\mathbf{c}, r)$ is denoted by $\mathcal{P}_{\mathbf{c},h}(P)$. In [CAGM97], the authors prove that any measure \tilde{P}^* realizing the infimum in the definition of $V_{k,h}(P)$ is an element of $\mathcal{P}_{\mathbf{c}^*,h}(P)$, for a codebook \mathbf{c}^* which minimizes $\mathbf{c} \mapsto V_h^P(\mathbf{c})$, where the *h -trimmed variation about \mathbf{c}* , $V_h^P(\mathbf{c})$ is defined for all $P_{\mathbf{c},h} \in \mathcal{P}_{\mathbf{c},h}(P)$ by

$$V_h^P(\mathbf{c}) = \inf_{\tilde{P} \in \mathcal{P}_{h+}(P)} V^{\tilde{P}}(\mathbf{c}) = V^{P_{\mathbf{c},h}}(\mathbf{c}).$$

In addition, see [CAGM97, Proposition 2.3], the trimmed k -variation coincides with the h -trimmed variation about \mathbf{c}^* , that is $V_h^P(\mathbf{c}^*) = V_{k,h}(P)$. Concretely, sets minimizing the loss are unions of balls centered at the elements of \mathbf{c}^* .

Such a method is particularly well adapted when the data are sampled from a mixture of Gaussian distributions, since in this case, sub-level sets of the density are union of Euclidean balls.

In [BMDG05], the authors propose a variation of k -means based on the Bregman divergence as an alternative to the squared Euclidean distance. Given some convex subset Ω of \mathbb{R}^d , and some strictly convex \mathcal{C}^1 function ϕ defined on Ω , the *Bregman divergence* to a point $y \in \Omega$ is defined for all $x \in \Omega$ by

$$d_\phi(x, y) = \phi(x) - \phi(y) - \langle \nabla_y \phi, x - y \rangle,$$

where $\nabla_y \phi$ denotes the gradient of ϕ at y .

The notion of Bregman divergence is particularly suited to mixtures of distributions from exponential families, such as for instance Poisson, binomial or even Gamma distributions. To each family of distributions, a function ϕ is naturally assigned. Moreover, the authors confirm, through experiments, that clustering data with the corresponding ϕ -Bregman divergence provides a better clustering. The quality of a clustering method will be measured in terms of normalized mutual information, introduced in [SJ02].

Another advantage of the use of Bregman divergences for clustering data comes from the possibility to adapt the Lloyd algorithm to compute a local minimizer of the loss. The adaptation is possible since for every strictly convex function ϕ , the expectation of P , Pu minimizes $c \mapsto Pd_\phi(u, c)$ among all elements $c \in \mathbb{R}^d$. Such a property is not satisfied for the function $x \mapsto \phi(\|x - c\|)$ used in [CAGM97]; the authors had to use an annealing based algorithm.

3.1.2 Contribution

In this Section, we propose a generalization of k -means to data sampled from a distribution P possibly corrupted by noise, where the mixtures of distributions from exponential families is of particular interest. The method we propose is an adaptation of [CAGM97] in the sense that a proportion $1 - h$ of the data will be wiped out. It might also be seen as an adaptation of [BMDG05] in the sense that the square of the Euclidean distance will be replaced by the Bregman divergence associated to some convex function ϕ .

Given $\mathbb{X}_n = \{X_1, X_2, \dots, X_n\}$, an n -sample from P , an *empirical optimal codebook* is an optimal codebook for the *empirical measure* $P_n = \frac{1}{n} \sum_{i=1}^n \delta_{X_i}$, the uniform measure on the set \mathbb{X}_n . We prove that empirical optimal codebooks converge a.e. to the optimal codebook under some assumptions, and that the loss associated to the empirical optimal codebook will converge to the loss associated to the optimal codebook (for P), at a rate $\frac{1}{\sqrt{n}}$. The proofs will rely on measure concentration inequalities such as the bounded difference inequality, together with the Dudley entropy integral and the metric entropy, related to the notion of shattering dimension.

We also make the link with the notion of distance-to-measure. Indeed, the Bregman h -trimmed variation about a codebook $\mathbf{c} = (c_1, c_2, \dots, c_k)$ coincides with the DTM when the Euclidean distance is replaced with the Bregman divergence, and the point x is replaced with \mathbf{c} and the ball centered at x with P -mass h is replaced with a union of balls centered at the c_i s with fixed radius and P -mass h . Then, the optimal codebooks coincide with the minima of this notion close to the DTM, that we will call Bregman divergence-to-measure.

In addition, we provide a Lloyd-type algorithm that is proved to converge to a local minimum for the empirical loss. We also furnish a heuristic for the selection of the trimming parameter h from data. The algorithm together with the heuristic are exploited on some illustrative examples where data are sampled according to a mixture of distributions from various exponential families, with clutter noise.

Before introducing a framework for trimmed clustering with the Bregman divergence, we introduce some additional notation. For $k \in \mathbb{N}^*$ and any space \mathcal{A} , $\mathcal{A}^{(k)}$ stands for $\{\mathbf{c} = (c_1, c_2, \dots, c_k) \mid \forall i \in \llbracket 1, k \rrbracket, c_i \in \mathcal{A}\}$, where two elements are identified whenever they are equal up to a permutation of the coordinates. We equip $\mathcal{A}^{(k)}$ with the metric $\text{dist}(\mathbf{c}, \mathbf{c}') = \inf_{\sigma \in \Sigma_k} \max_{i \in \llbracket 1, k \rrbracket} \|c_i - c'_{\sigma_i}\|$, with Σ_k the set of all permutations of $\llbracket 1, k \rrbracket$. Also, for any subset A of \mathbb{R}^d , \overline{A} stands for its closure, A° for its interior and $A^c = \mathbb{R}^d \setminus A$ its complementary set in \mathbb{R}^d .

3.1.3 A framework for trimmed clustering with Bregman divergences

A natural setting adapted from the trimmed k -means setting

Definition 3.1. Let ϕ be a strictly convex \mathcal{C}^1 real-valued function defined on a convex set $\Omega \subset \mathbb{R}^d$, the *Bregman divergence* d_ϕ is defined for all $x, y \in \Omega$ by

$$d_\phi(x, y) = \phi(x) - \phi(y) - \langle \nabla_y \phi, x - y \rangle.$$

We generalise this notion to every codebook $\mathbf{c} = (c_1, c_2, \dots, c_k) \in \Omega^{(k)}$ by:

$$d_\phi(x, \mathbf{c}) = \min_{i \in \llbracket 1, k \rrbracket} d_\phi(x, c_i).$$

Note that by taking $\phi : x \mapsto \|x\|^2$, $d_\phi(x, y) = \|x - y\|^2$. Also, since ϕ is strictly convex, for all $x, y \in \mathbb{R}^d$, $d_\phi(x, y)$ is non-negative and equal to zero if and only if $x = y$; see [Roc70, Theorem 25.1].

Lemma 3.2 (Proposition 2 in [BMDG05]). *The map $c \in \Omega \mapsto Pd_\phi(u, c)$ has a unique minimizer given by $c = Pu$.*

Proof. For all $c \in \Omega$, $Pd_\phi(u, c) = Pd_\phi(u, Pu) + d_\phi(Pu, c)$. □

Note that this minimizer belongs to Ω which is a convex set.

Definition 3.3. The Bregman ball $\mathcal{B}_\phi(\mathbf{c}, r)$ centered at $\mathbf{c} \in \Omega^{(k)}$, with radius $r > 0$ is defined by

$$\mathcal{B}_\phi(\mathbf{c}, r) = \left\{ x \in \Omega \mid \sqrt{d_\phi(x, \mathbf{c})} < r \right\}.$$

Note the presence of a square root on the Bregman divergence in the definition of Bregman ball. It comes from the fact that the Bregman divergence behaves like a square; see for instance the particular case where ϕ equals the squared Euclidean distance or the more general case where ϕ is *strongly convex* (i.e. s.t. $d_\phi(x, y) \geq \frac{\mu}{2} \|x - y\|^2$ for some $\mu > 0$ and all $x, y \in \Omega$).

For all $h \in (0, 1)$, we will use the notation $r_{\phi, h}(\mathbf{c})$ for the largest radius $r \geq 0$ such that

$$P(\mathcal{B}(\mathbf{c}, r)) \leq h \leq P(\overline{\mathcal{B}}(\mathbf{c}, r)). \quad (3.1)$$

Also, denote by $\mathcal{P}_{\phi, \mathbf{c}, h}(P)$, the set of all probability measures $P_{\mathbf{c}, h} \in \mathcal{P}_h(P)$ such that $P_{\mathbf{c}, h}$ is supported on $\overline{\mathcal{B}}_\phi(\mathbf{c}, r)$ and $hP_{\mathbf{c}, h}$ coincides with P on $\mathcal{B}_\phi(\mathbf{c}, r)$.

Similarly to [CAGM97], we can define notions of trimmed variation:

Definition 3.4. For any $P \in \mathcal{P}(\Omega)$, the *variation about $\mathbf{c} \in \Omega^{(k)}$ given P* , $V_\phi^P(\mathbf{c})$ is defined by:

$$V_\phi^P(\mathbf{c}) = Pd_\phi(u, \mathbf{c}).$$

Given $h \in (0, 1)$, the *h -trimmed variation about $\mathbf{c} \in \Omega^{(k)}$* , $V_{\phi, h}^P(\mathbf{c})$ is defined for any $P_{\mathbf{c}, h}$ in $\mathcal{P}_{\phi, \mathbf{c}, h}(P)$ by:

$$V_{\phi, h}^P(\mathbf{c}) = P_{\mathbf{c}, h}d_\phi(u, \mathbf{c}).$$

The *k -variation given P* , $V_{\phi, k}^P$ is defined by:

$$V_{\phi, k}^P = \inf_{\mathbf{c} \in \Omega^{(k)}} V_\phi^P(\mathbf{c}).$$

The *h -trimmed k -variation*, $V_{\phi, k, h}(P)$ is defined by:

$$V_{\phi, k, h}(P) = \inf_{\tilde{P} \in \mathcal{P}_{h^+}(P)} V_{\phi, k}^{\tilde{P}}.$$

While the variation about \mathbf{c} corresponds to the P -mean of the Bregman divergence to \mathbf{c} on the whole space Ω , the h -trimmed variation about \mathbf{c} corresponds to the P -mean of the Bregman divergence to \mathbf{c} on the Bregman ball centered at \mathbf{c} with P -mass h .

The k -variation given P stands for the variation about an optimal codebook, that is, a minimizer of the variation. Such a codebook is optimal in the sense of a clustering with Bregman divergence, see [BMDG05]. Finally, the h -trimmed k -variation corresponds to the best clustering with the Bregman divergence, when we allow ourselves to remove a proportion of mass at most $1 - h$ of the measure P .

When fixing a codebook $\mathbf{c} \in \Omega^{(k)}$, the best h -trimming procedure consists in keeping the closest points to \mathbf{c} in terms of the Bregman divergence to \mathbf{c} , that is, the points within the Bregman ball centered at \mathbf{c} with P -mass h . In this sense, we will prove through the two following lemmas, that the h -trimmed variation about $\mathbf{c} \in \Omega^{(k)}$ can be rewritten as $\inf_{\tilde{P} \in \mathcal{P}_{h^+}(P)} \tilde{P}d_\phi(u, \mathbf{c})$.

Lemma 3.5. For all $\mathbf{c} \in \Omega^{(k)}$, $h \in (0, 1]$ and $\tilde{P} \in \mathcal{P}_h(P)$,

$$V_{\phi,h}^P(\mathbf{c}) \leq \tilde{P}d_\phi(u, \mathbf{c}).$$

Equality holds if and only if $\tilde{P} \in \mathcal{P}_{\phi,\mathbf{c},h}(P)$.

This lemma is a straightforward generalisation of results in [CAGM97, Lemma 2.1] or [Gor91], its proof can be found in Section 3.3.1.

Lemma 3.6. Let $0 < h < h' < 1$. Then, for all $P \in \mathcal{P}(\Omega)$ and $\mathbf{c} \in \Omega^{(k)}$, we have:

$$V_{\phi,h}^P(\mathbf{c}) \leq V_{\phi,h'}^P(\mathbf{c}).$$

Moreover, equality holds if and only if $P(\mathcal{B}_\phi(\mathbf{c}, r_{\phi,h'}(\mathbf{c}))) = 0$.

This lemma is a straightforward generalisation of results in [CAGM97, Lemma 2.2], see also [Gor91]. Its proof can be found in Section 3.3.1.

Minimizing the h -trimmed k -variation boils down to finding the codebook that minimizes the h -trimmed variation, as described below.

Proposition 3.7.

$$V_{\phi,k,h}(P) = \inf_{\mathbf{c} \in \Omega^{(k)}} V_{\phi,h}^P(\mathbf{c}).$$

Proof. This proposition is a straightforward generalisation of Proposition 2.3 in [CAGM97]. According to Lemma 3.6 and Lemma 3.5, $V_{\phi,h}^P(\mathbf{c}) = \inf_{\tilde{P} \in \mathcal{P}_{h^+}(P)} \tilde{P}d_\phi(u, \mathbf{c})$. The result follows from Definition 3.4. \square

Thus, the target quantity is any minimizer \mathbf{c} of the h -trimmed variation. Then, it will lead to a pertinent clustering of the data.

Definition 3.8. A h -trimmed k -optimal codebook is any element \mathbf{c}^* in $\arg \min_{\mathbf{c} \in \Omega^{(k)}} V_{\phi,h}^P(\mathbf{c})$.

The performance of a clustering might be measured in terms of the optimal risk $R_{k,h}^* = hV_{\phi,h}^P(\mathbf{c}^*) = hV_{\phi,k,h}(P)$.

Such a trimmed k -optimal codebook exists under some assumptions on P and ϕ , as follows.

Theorem 3.9. Assume that ϕ is \mathcal{C}^2 and strictly convex and $F_0 = \overline{\text{Conv}(\text{Supp}(P))} \subset \overset{\circ}{\Omega}$. Then, the set $\arg \min_{\mathbf{c} \in \Omega^{(k)}} V_{\phi,h}^P(\mathbf{c})$ is not empty.

The proof of Theorem 3.9 is to be found in Section 3.3.1.

Trimmed k -optimal codebook and Bregman-Voronoi measures

Any measure P can be decomposed into a set of $k \in \mathbb{N}^*$ Bregman-Voronoi measures, supported on the Bregman-Voronoi cells $V_i = \{x \in \Omega \mid \forall j \neq i, d_\phi(x, c_i) \leq d_\phi(x, c_j)\}$ for $i \in \llbracket 1, k \rrbracket$. Bregman-Voronoi cells were studied in [BNN10] for instance.

Definition 3.10. A set of Bregman-Voronoi measures associated to a distribution $P \in \mathcal{P}_1(\mathbb{R}^d)$ and $\mathbf{c} = (c_1, c_2, \dots, c_k) \in \Omega^{(k)}$ is a set $\{P^{BV,i,\mathbf{c}}\}_{i \in \llbracket 1, k \rrbracket}$ of k positive sub-measures of P such that $\sum_{i=1}^k P^{BV,i,\mathbf{c}} = P$ and $\text{Supp}(P^{BV,i,\mathbf{c}}) \subset V_i$. We denote by $m(P^{BV,i,\mathbf{c}}) = \frac{P^{BV,i,\mathbf{c}}(u)}{P^{BV,i,\mathbf{c}}(\mathbb{R}^d)}$ the expectation of $P^{BV,i,\mathbf{c}}$, and set $m(P^{BV,i,\mathbf{c}}) = c_0$ when $P^{BV,i,\mathbf{c}}(\mathbb{R}^d) = 0$, for some arbitrary element $c_0 \in \text{Supp}(P)$.

For $h \in (0, 1)$ we define the operator T_h (or T) on $\Omega^{(k)}$ for any choice of measure $P_{\mathbf{c},h} \in \mathcal{P}_{\mathbf{c},h}(P)$ as follows:

$$T : \begin{cases} \Omega^{(k)} & \rightarrow \Omega^{(k)} \\ \mathbf{c} = (c_1, \dots, c_k) & \rightarrow \left(m \left(P_{\mathbf{c},h}^{BV,i,\mathbf{c}} \right) \right)_{i=1,\dots,k} \end{cases}.$$

When P puts no mass on the boundaries of the cells $V_{i,h} = V_i \cap \text{Supp}(P_{\mathbf{c},h})$, we may write $T(\mathbf{c}) = \frac{P \mathbb{1}_{V_{i,h}}(u)}{P(V_{i,h})}$.

The h -trimmed variation is non-increasing with T . As a consequence, the following property is satisfied for optimal codebooks.

Proposition 3.11. *For every $h \in (0, 1)$ and $\mathbf{c} \in \Omega$, $V_{\phi,h}^P(T(\mathbf{c})) \leq V_{\phi,h}^P(\mathbf{c})$. Moreover, if \mathbf{c}^* is a trimmed k -optimal codebook of P for the mass parameter h . Then, for every $P_{\mathbf{c}^*,h} \in \mathcal{P}_{\phi,\mathbf{c}^*,h}(P)$,*

$$\forall i \in \llbracket 1, k \rrbracket, c_i^* = T_h(\mathbf{c}^*)_i = m(P_{\mathbf{c}^*,h}^{BV,i,\mathbf{c}^*}) \text{ if } P_{\mathbf{c}^*,h}^{BV,i,\mathbf{c}^*}(\Omega) \neq 0,$$

where $\left\{ P_{\mathbf{c}^*,h}^{BV,i,\mathbf{c}^*} \right\}_{i \in \llbracket 1, k \rrbracket}$ is a set of Bregman-Voronoi measures associated to $P_{\mathbf{c}^*,h}$.

The proof of Proposition 3.11, which is a generalisation of Proposition 3.3 in [CAGM97] is to be found in Section 3.3.1.

It is possible to derive bounds for the difference of the h -trimmed variations at two points, as follows.

Lemma 3.12. *Let $P \in \mathcal{P}_1(\Omega)$, $h \in (0, 1)$ and $\mathbf{c}, \mathbf{c}' \in \Omega^{(k)}$. For some element $P_{\mathbf{c},h}$ in $\mathcal{P}_{\phi,\mathbf{c},h}(P)$, set $\left\{ P_{\mathbf{c},h}^{BV,i,\mathbf{c}} \right\}_{i \in \llbracket 1, k \rrbracket}$ a set of Bregman-Voronoi measures associated to $P_{\mathbf{c},h}$. Then, the two following upper bounds are satisfied*

$$\begin{aligned} V_{\phi,h}^P(\mathbf{c}') - V_{\phi,h}^P(\mathbf{c}) &\leq \sum_{i=1}^k P_{\mathbf{c},h}^{BV,i,\mathbf{c}}(\Omega) \left(\langle \nabla_{c_i} \phi - \nabla_{c'_i} \phi, m \left(P_{\mathbf{c},h}^{BV,i,\mathbf{c}} \right) - c'_i \rangle - d_\phi(c'_i, c_i) \right), \\ V_{\phi,h}^P(\mathbf{c}') - V_{\phi,h}^P(\mathbf{c}) &\leq \sum_{i=1}^k P_{\mathbf{c},h}^{BV,i,\mathbf{c}}(\Omega) \left(\langle \nabla_{c_i} \phi - \nabla_{c'_i} \phi, m \left(P_{\mathbf{c},h}^{BV,i,\mathbf{c}} \right) - c_i \rangle + d_\phi(c_i, c'_i) \right). \end{aligned}$$

In particular,

$$V_{\phi,h}^P(\mathbf{c}') - V_{\phi,h}^P(\mathbf{c}) \leq \sum_{i=1}^k P_{\mathbf{c},h}^{BV,i,\mathbf{c}}(\Omega) \left\| \nabla_{c'_i} \phi - \nabla_{c_i} \phi \right\| \left\| m \left(P_{\mathbf{c},h}^{BV,i,\mathbf{c}} \right) - c'_i \right\|.$$

The proof of Lemma 3.12 is to be found in Section 3.3.1.

It follows that the h -trimmed variation associated to $P \in \mathcal{P}_1(\Omega)$ is stable with respect to the Bregman divergence in the neighborhood of any trimmed k -optimal codebook of P in the following sense.

Proposition 3.13. *Let $P \in \mathcal{P}_1(\Omega)$, $h \in [0, 1]$, \mathbf{c}^* a trimmed k -optimal codebook for P , $\mathbf{c} \in \Omega^{(k)}$ and $\left\{ P_{\mathbf{c}^*,h}^{BV,i,\mathbf{c}^*} \right\}_{i \in \llbracket 1, k \rrbracket}$ the Bregman-Voronoi measures associated to some $P_{\mathbf{c}^*,h} \in \mathcal{P}_{\phi,\mathbf{c}^*,h}(P)$. Then,*

$$0 \leq V_{\phi,h}^P(\mathbf{c}) - V_{\phi,h}^P(\mathbf{c}^*) \leq \sum_{i=1}^k P_{\mathbf{c}^*,h}^{BV,i,\mathbf{c}^*}(\Omega) d_\phi(c_i^*, c_i).$$

Proof. From Proposition 3.11, since \mathbf{c}^* is optimal, we get that $c_i^* = m \left(P_{\mathbf{c}^*,h}^{BV,i,\mathbf{c}^*} \right)$ for all i such that $P_{\mathbf{c}^*,h}^{BV,i,\mathbf{c}^*}(\Omega) \neq 0$. We conclude with Lemma 3.12. \square

Convergence of the empirical optimal codebook to the optimal codebook

Given an n -sample from P , a natural approximation of the trimmed k -optimal codebook \mathbf{c}^* is given by an empirical trimmed k -optimal codebook, i.e. an element $\hat{\mathbf{c}}_n$ of $\arg \min_{\mathbf{c} \in \Omega^{(k)}} V_{\phi,h}^{P_n}(\mathbf{c})$. This strategy makes more sense when \mathbf{c}^* is unique as a minimizer of the variation given P . In this case, under some continuity assumptions on P , $\hat{\mathbf{c}}_n$ is proved to converge to \mathbf{c}^* almost everywhere for the distance dist , when n goes to infinity. Indeed, it follows from a straightforward generalisation of Theorem 3.4 in [CAGM97]:

Theorem 3.14. *If P is continuous, $P\|u\|^p < \infty$ for some $p > 2$, ϕ is \mathcal{C}_2 on Ω , $F_0 = \overline{\text{Conv}(\text{Supp}(P))} \subset \Omega^\circ$ and \mathbf{c}^* is the unique minimizer of $V_{\phi,h}^P$, then:*

$$\lim_{n \rightarrow +\infty} \text{dist}(\hat{\mathbf{c}}_n, \mathbf{c}^*) = 0 \text{ a.e.}$$

where $\text{dist}(\mathbf{c}, \mathbf{c}') = \min_{\sigma \in \Sigma_k} \max_{i \in \llbracket 1, k \rrbracket} |c_i - c'_{\sigma(i)}|$ where Σ_k denotes the set of all permutations of $\llbracket 1, k \rrbracket$ and

$$\lim_{n \rightarrow +\infty} V_{\phi,k,h}(P_n) = V_{\phi,k,h}(P) \text{ a.e.}$$

Proof. The proof of Theorem 3.14 is to be found in Section 3.3.1. \square

In particular, the convergence of the loss in the empirical codebook to the optimal loss holds at a rate $\frac{1}{\sqrt{n}}$.

Theorem 3.15. *Assume that $P\|u\|^p < \infty$. Then, for n large enough, with probability greater than $1 - n^{-\frac{p}{2}} - 2e^{-x}$, we have*

$$V_{\phi,h}^P(\hat{\mathbf{c}}_n) - V_{\phi,k,h}(P) \leq \frac{C_P}{h\sqrt{n}}(1 + \sqrt{x}).$$

The proof of Theorem 3.15 is to be found in Section 3.3.1.

Corollary 3.16. *If in addition there exist $c_0 \in \Omega^\circ$ and ψ a convex function such that*

$$\sup_{c \in \mathcal{B}(c_0, t) \cap F_0} \|\nabla_c \phi\| \leq \psi(t),$$

with $P\|u\|^2 \psi^2\left(\frac{k\|u\|}{h}\right) < +\infty$, $P\|u\|^2 < +\infty$ and $P\psi^2\left(\frac{k\|u\|}{h}\right) < +\infty$, then

$$\mathbb{E} \left(V_{\phi,h}^P(\hat{\mathbf{c}}_n) - V_{\phi,k,h}(P) \right) \leq \frac{C_P}{h\sqrt{n}}.$$

The proof of Corollary 3.16 is to be found in Section 3.3.1.

Optimal codebooks and the distance to measure

The h -trimmed variation corresponds to a slightly modified version of the distance-to-measure, as introduced by [CCSM11]. We will introduce the notion of Bregman divergence-to-measure that corresponds to this slight modification, as follows.

Definition 3.17. Given P the Bregman divergence to the measure P is defined for every codebook $\mathbf{c} \in \Omega$ by

$$d_{\phi,P,h}^2(\mathbf{c}) = \inf_{\tilde{P} \in \mathcal{P}_h(P)} \tilde{P}d_{\phi}(u, \mathbf{c}),$$

or equivalently, according to Lemma 3.5, by

$$d_{\phi,P,h}^2(\mathbf{c}) = P_{\mathbf{c},h}d_{\phi}(u, \mathbf{c})$$

for every $P_{\mathbf{c},h} \in \mathcal{P}_{\phi,\mathbf{c},h}(P)$ as defined in Section 3.1.3.

According to Definition 3.4, Lemma 3.5, Lemma 3.6 and Proposition 3.7, the notion of Bregman divergence-to-measure coincides with the notion of h -trimmed variation. In particular, optimal codebooks correspond to minimisers of the Bregman divergence to the measure P , and are approximated by any minimiser $\hat{\mathbf{c}}_n$ of the Bregman divergence to the measure P_n .

The arguments in [CMM16] generalise to the Bregman divergence to derive deviation bounds for $\sup_{\mathbf{c}} |d_{\phi,P,h}^2(\mathbf{c}) - d_{\phi,P_n,h}^2(\mathbf{c})|$. It is thus possible to get faster rates of convergence for h small, as a function of the continuity modulus of the quantile of the distribution of $d_{\phi}^2(X, \mathbf{c})$ for $X \sim P$ and \mathbf{c} in a compact set, and the covering number of $\mathbf{c} \mapsto d_{\phi,P,h}^2(\mathbf{c})$. Such a remark might be useful to bound $d_{\phi,P,h}^2(\hat{\mathbf{c}}_n) - d_{\phi,P,h}^2(\mathbf{c}^*)$. Nonetheless, in the context of trimmed clustering, the proportion of signal should be close to h , which should not go to zero. Thus, in this context, the rates of convergence and deviation bounds obtained should not be better than the ones obtained in Theorem 3.15 and Corollary 3.16.

In [CCSM11], the authors prove that the distance-to-measure is 1-Lipschitz. Such a property is not necessarily satisfied by the Bregman divergence-to-measure, but at least it remains continuous over compact sets, as follows.

Proposition 3.18. Let P be a probability measure such that $P\|u\| < \infty$, $\text{Supp}(P) \subset \mathcal{B}(0, K)$ and $\mathbf{c}, \mathbf{c}' \in \mathcal{K}^{(k)}$ for \mathcal{K} some convex compact subset of $\Omega \cap \mathcal{B}(0, K')$ for some $K' > 0$. Then,

$$|d_{\phi,P,h}^2(\mathbf{c}) - d_{\phi,P,h}^2(\mathbf{c}')| \leq (K + K') \omega_{\nabla\phi,\mathcal{K}}(\text{dist}(\mathbf{c}, \mathbf{c}')),$$

where $\omega_{\nabla\phi,\mathcal{K}}$ is the continuity modulus of the function $c \in \mathcal{K} \mapsto \nabla\phi_c$, that is, is such that

$$\forall \epsilon > 0, \omega_{\nabla\phi,\mathcal{K}}(\epsilon) = \sup \{ \|\nabla\phi_c - \nabla\phi_{c'}\| \mid c, c' \in \mathcal{K}, \|c - c'\| \leq \epsilon \}.$$

Proof. It comes from Lemma 3.12 and the fact that $\|m(P_{\mathbf{c},h}^{BV,i,\mathbf{c}})\| \leq K$ since $\text{Supp}(P_{\mathbf{c},h}^{BV,i,\mathbf{c}}) \subset \text{Supp}(P)$ when $P_{\mathbf{c},h} \in \mathcal{P}_{\phi,\mathbf{c},h}(P)$. \square

This proposition can be extended to general measures.

Proposition 3.19. Let P be an inner-regular probability measure such that $P\|u\| < \infty$ and $\text{Supp}(P) \subset \Omega$. Set $\mathbf{c}, \mathbf{c}' \in \mathcal{K}^{(k)}$, where \mathcal{K} is a convex compact subset of $\Omega^\circ \cap \mathcal{B}(0, K')$ for some $K' > 0$, then for every $h \in (0, 1)$,

$$|d_{\phi,P,h}^2(\mathbf{c}) - d_{\phi,P,h}^2(\mathbf{c}')| \leq L_{\phi,P,\mathcal{K}} \omega_{\nabla\phi,\mathcal{K}}(\text{dist}(\mathbf{c}, \mathbf{c}')).$$

for some finite Lipschitz constant $L_{\phi,P,\mathcal{K}}$ depending on ϕ , P , \mathcal{K} and h .

The proof of Proposition 3.19 is to be found in Section 3.3.1.

Note that since ϕ is supposed to be \mathcal{C}^1 , $c \mapsto \nabla\phi_c$ is uniformly continuous on the compact set \mathcal{K} . Thus the continuity modulus is finite for all ϵ and goes to zero when ϵ goes to zero. For instance, choosing $\phi(c) = \|c\|^2$ leads to $\omega_{\nabla\phi,\mathcal{K}}(\epsilon) = 2\epsilon$ for any compact set \mathcal{K} .

In particular, $\mathbf{c} \mapsto d_{\phi,P,h}^2(\mathbf{c})$ is continuous whenever ϕ is \mathcal{C}^1 .

3.1.4 Algorithm

The following algorithm is inspired by Lloyd's algorithm, but is also a generalization of the Bregman Hard Clustering for uniform finite-supported measures; see [BMDG05, Algorithm 1]. We assume that the mass parameter $h = \frac{q}{n}$ for some positive integer q . Set \mathcal{C}_i the Bregman-Voronoi cell associated to c_i . We use the notation $|\mathcal{C}_i|$ for the cardinal of $\mathcal{C}_i \cap \mathbb{X}_n$.

Algorithm 3.1: Bregman trimmed k-means algorithm

```

Input :  $\mathbb{X}_n$  an  $n$ -sample from  $P$ ,  $q$  and  $k$  ;
# Initialization
Sample  $c_1, c_2, \dots, c_k$  from  $\mathbb{X}_n$  without replacement ;
while the  $c_i$ s vary make the following two steps :
  # Decomposition into Bregman - Voronoi cells .
  for  $i$  in  $1..k$ :
     $\mathcal{C}_i \leftarrow \emptyset$ 
  for  $j$  in  $1..n$ :
    Add  $X_j$  to the  $\mathcal{C}_i$  (for  $i$  as small as possible) satisfying
     $d_\phi(X_j, c_i) \leq d_\phi(X_j, c_l), \forall l \neq i$ ;
     $d_{min}(j) = \min_{i=1..k} d_\phi(X_j, c_i)$ ;
  # Trimming stage
  Sort  $d_{min}$ ;
  Set  $I$  the set of indices associated to the  $n - q$  largest elements of  $d_{min}$ ;
   $cost = \sum_{i \in [1, n] \setminus I} (d_{min}[i])$ ;
  for  $j$  in  $I$ :
    Remove  $X_j$  from its associated cell  $\mathcal{C}_i$ ;
  # Computation of the new centers and weights .
  for  $i$  in  $1..k$ :
     $c_i = \frac{1}{|\mathcal{C}_i|} \sum_{X \in \mathcal{C}_i} X$ ;
Output :  $(c_1, c_2, \dots, c_k), (\mathcal{C}_1, \mathcal{C}_2, \dots, \mathcal{C}_k)$  and cost

```

Proposition 3.20. *The algorithm 3.1 converges to a local minimum of the function $V_{\phi, h_n}^{P_n}$ with $h_n = \frac{q}{n}$.*

The proof of Proposition 3.20 is an adaptation of the proof of Proposition 2 in [BMDG05]. It is a straightforward consequence of Proposition 3.11 since at each step, c_i is replaced with $T_n(c_i)$, where T_n is the operator T associated with the measure P_n .

3.1.5 Experiments

In this part, we apply Algorithm 3.1 to mixtures of distributions belonging to some exponential family. As noted in [BMDG05], we can associate to some exponential families of distribution a Bregman divergence, by Legendre duality. Such a Bregman divergence is more adapted for the clustering than other divergences, as shown by the following experiments. First we recall that an exponential family associated to a proper closed convex function ψ defined on an open parameter space $\Theta \subset \mathbb{R}^d$ is a family of distributions $\mathcal{F}_\psi = \{P_{\psi, \theta} \mid \theta \in \Theta\}$ satisfying the following properties. For all $\theta \in \Theta$, $P_{\psi, \theta}$ is defined on a space \mathbb{R}^d and is absolutely continuous with respect to some distribution P_0 with Radon-Nikodym density $p_{\psi, \theta}$ defined for all $x \in \Omega$ by:

$$p_{\psi, \theta}(x) = \exp(\langle x, \theta \rangle - \psi(\theta)).$$

Since $P_{\psi, \theta}$ is a probability distribution, the differentiation of the $L_1(P_0)$ norm of its density yields $\mu(\theta) = \nabla_\theta \psi$, where $\mu(\theta)$ is the expectation of $p_{\psi, \theta}$. We define $\phi(\mu) = \sup_{\theta \in \Theta} \{\langle \mu, \theta \rangle - \psi(\theta)\}$.

Distribution	$p(x, \theta)$	θ	$\psi(\theta)$
Gaussian	$\frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(x-a)^2}{2\sigma^2}\right)$	$\frac{a}{\sigma^2}$	$\frac{\sigma^2}{2}\theta^2$
Poisson	$\frac{\lambda^x \exp(-\lambda)}{x!}$	$\log(\lambda)$	$\exp(\theta)$
Binomial	$\frac{N!}{x!(N-x)!} q^x (1-q)^{N-x}$	$\log\left(\frac{q}{1-q}\right)$	$N \log(1 + \exp(\theta))$
Gamma	$\frac{x^{k-1} \exp(-\frac{x}{b})}{\Gamma(k)b^k}$	$-\frac{k}{\mu}$	$k \log\left(-\frac{1}{\theta}\right)$

Distribution	μ	$\phi(\mu)$	$d_\phi(x, \mu)$
Gaussian	a	$\frac{1}{2\sigma^2}\mu^2$	$\frac{1}{2\sigma^2}(x - \mu)^2$
Poisson	λ	$\mu \log(\mu) - \mu$	$x \log\left(\frac{x}{\mu}\right) - (x - \mu)$
Binomial	Nq	$\mu \log\left(\frac{\mu}{N}\right) + (N - \mu) \log\left(\frac{N-\mu}{N}\right)$	$x \log\left(\frac{x}{\mu}\right) + (N - x) \log\left(\frac{N-x}{N-\mu}\right)$
Gamma	kb	$-k + k \log\left(\frac{k}{\mu}\right)$	$\frac{k}{\mu} \left(\mu \log\left(\frac{\mu}{x}\right) + x - \mu\right)$

Figure 3.1: Exponential families and associated Bregman divergence

According to [BMDG05, Theorem 2] extracted from [Roc70], by Legendre duality, for all μ where ϕ is defined, we get:

$$\phi(\mu) = \langle \theta(\mu), \mu \rangle - \psi(\theta(\mu)),$$

with $\theta(\mu) = \nabla_\mu \phi$. Moreover, according to [BMDG05, Theorem 4], the density of $P_{\psi, \theta}$ with respect to P_0 can be rewritten with the Bregman divergence associated to ϕ as follows:

$$p_{\psi, \theta}(x) = \exp(-d_\phi(x, \mu) + \phi(x)).$$

We consider the example of the family of Gamma distributions.

Example 3.21. The density of the Gamma distribution with respect to the Lebesgue measure on $]0, +\infty[$ is given by:

$$x \mapsto \frac{x^{k-1} \exp\left(-\frac{x}{b}\right)}{\Gamma(k)b^k} = \exp(\langle x, \theta \rangle - \psi(\theta) + \mathbf{c}(x)),$$

with $\theta = \frac{-1}{b}$ and $\psi(\theta) = k \log\left(\frac{-1}{\theta}\right)$. Here, Γ denotes the Gamma function. Then $\mu(\theta) = \nabla_\theta \psi = -\frac{k}{\theta}$. Thus, $\theta(\mu) = \frac{-k}{\mu}$.

Set $\phi(\mu) = \langle \mu, \theta(\mu) \rangle - \psi(\theta(\mu)) = -k + k \log\left(\frac{k}{\mu}\right)$. Then, $\nabla_\mu \phi = -\frac{k}{\mu}$. Thus,

$$d_\phi(x, \mu) = \phi(x) - \phi(\mu) - \langle x - \mu, \nabla_\mu \phi \rangle = \frac{k}{\mu} \left(\mu \log\left(\frac{\mu}{x}\right) + x - \mu\right).$$

We can verify that the density is really equal to $C(x) \exp(-d_\phi(x, \mu))$ for some nonnegative function C of x .

The experiments are made from Gaussian, Poisson, Binomial and Gamma distributions, but also for the Cauchy distribution. Figure 3.1 contains the densities together with the functions ϕ and the associated Bregman divergences. For more distributions, see [BMDG05, Table 2].

Appropriate choice of the Bregman divergence for mixtures

Before introducing noise, we redo the experiment of [BMDG05]. We generate 100-samples from different distributions which are various mixtures of Gaussian, Poisson, Binomial and Gamma distribution, but also from the Cauchy distribution. Each time there are three components of the mixtures with equal probability $\frac{1}{3}$. The means are set to 10, 20 and 40 (or the modes for the Cauchy distribution). For the Gaussian densities, we set the standard deviation to 5, for the Binomial distribution, the number of trials to 100 and for the Gamma distribution, the shape parameter to 40. The following figure shows the densities of the different mixtures when they do have density, or the diagram of probabilities for the discrete distributions.

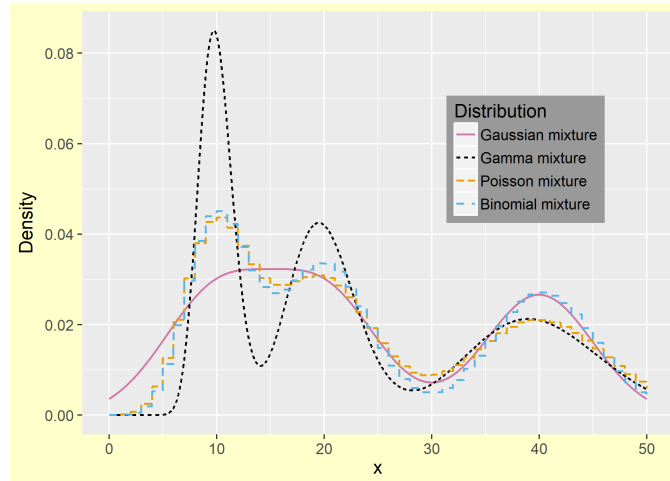


Figure 3.2: Distributions

For each distribution and each divergence, we sampled 100 100-samples. For each sample, we used algorithm 3.1 10 times and kept the best clustering, that is, the clustering for which the mean of d_{min} is minimal. Then we computed the normalized mutual information NMI , as defined in [SJ02] (see Definition 3.22) between the true clusters and the clusters obtained with algorithm 3.1.

Definition 3.22 ([SJ02]). The *normalized mutual information* is defined for two clusterings (a) and (b) by

$$NMI((a), (b)) = \frac{\sum_{h=1}^{k^{(a)}} \sum_{l=1}^{k^{(b)}} n_{h,l} \log \left(\frac{nn_{h,l}}{n_h^{(a)} n_l^{(b)}} \right)}{\sqrt{\sum_{h=1}^{k^{(a)}} n_h^{(a)} \log \left(\frac{n_h^{(a)}}{n} \right)} \sqrt{\sum_{l=1}^{k^{(b)}} n_l^{(b)} \log \left(\frac{n_l^{(b)}}{n} \right)}}.$$

We use the notation $k^{(a)}$ for the number of groups in the clustering (a) and $n_h^{(a)}$ the number of points assigned to the h -th group for the clustering (a) . Similarly, we associate $h^{(b)}$ and $n_l^{(b)}$ to the clustering (b) . Also, $n_{h,l}$ corresponds to the number of points that are both in the h -th group for clustering (a) and in the l -th group for clustering (b) .

In Figure 3.3, we derive confidence intervals for the NMI with level around 5 percent centred at the mean $mean(NMI)$ and with width $2 \times 1.96 \sqrt{\frac{var(NMI)}{100}}$.

As noted in [BMDG05], the highest renormalised mutual information is obtained when the Bregman divergence used is the one associated to the distribution.

Generative Model	$d_{Gaussian}$	$d_{Poisson}$	$d_{Binomial}$	d_{Gamma}
Gaussian	0.6857 ± 0.0126	0.6634 ± 0.0126	0.6754 ± 0.0131	0.4673 ± 0.0534
Poisson	0.6978 ± 0.0138	0.7418 ± 0.0123	0.7307 ± 0.0143	0.7082 ± 0.0175
Binomial	0.7888 ± 0.0124	0.8090 ± 0.0131	0.8132 ± 0.0117	0.7623 ± 0.0152
Gamma	0.8448 ± 0.0166	0.9083 ± 0.0111	0.8911 ± 0.0119	0.9177 ± 0.0102
Cauchy	0.7091 ± 0.0116	0.6941 ± 0.0126	0.7023 ± 0.0129	0.6761 ± 0.0109

Figure 3.3: Normalised mutual information

Trimming parameter selection

In this work, the interest resides in the existence of noise in the data. Thus, we made the same experiment but with a level of noise. First we propose a heuristic to select the trimming parameter q , that is, the number of points in the sample that are assigned to a cluster and not considered as noise in the procedure. We propose to proceed as follows.

Algorithm 3.2: Trimming parameter selection

```

Input :  $\mathbb{X}_n$  an  $n$ -sample from  $P$  and  $k$  ;
for  $q$  in  $1..n$ :
     $cost[q] \leftarrow \text{Algo 3.1}(\mathbb{X}_n, q, k)$  ;
Output : First jump of the curve  $q \mapsto cost[q]/q$ 
    
```

Morally, when the parameter q gets large enough, we start assigning outliers to clusters. These outliers will have large cost, thus there should be a jump in the slope.




Figure 3.4: Gaussian mixture




Figure 3.5: Poisson mixture




Figure 3.6: Binomial mixture




Figure 3.7: Gamma mixture

Figure 3.8: Trimming parameter q selection

Then in Figure 3.1.5 we have plotted the clustering associated to the choice of $q = 110$ for the Gaussian mixture, $q = 106$ for the Poisson mixture, $q = 109$ for the Binomial mixture and $q = 104$ for the Gamma mixture, the parameters that are given by the heuristic. We selected the parameter $q = 105$ for the Cauchy distribution.




Figure 3.9: Gaussian mixture




Figure 3.10: Poisson mixture




Figure 3.11: Binomial mixture




Figure 3.12: Gamma mixture

Figure 3.13: Clustering associated to the selected parameter q




Figure 3.14: Cauchy mixture




Figure 3.15: Cauchy mixture

We also made the experiment with the Cauchy distribution, with the square of the Euclidean norm, that is, the Bregman divergence associated with the Gaussian distribution. The selecting of the trimming parameter works well and the associated clustering is also good. Thus, the method adapts to mixture models that are not necessarily from distributions in exponential families. Note that the square Euclidean distance is particularly well adapted since unlike Poisson, Binomial and Gamma distributions, Gaussian and Cauchy distributions are both symmetrical.

Appropriate choice of the Bregman divergence for mixtures with noise

We do the same experiment as in Section 3.1.5 on 100 120-samples with each time 100 points from the mixture distribution and 20 outliers which are uniformly sampled on $[0, 200]$, like in Section 3.1.5. Most of the time, with this number of points, our heuristic leads to a selection of q around 103. In the following experiments, we choose $q = 103$ and proceed as in Section 3.1.5. The points which are noise are attributed to the same cluster. We then compute the renormalized mutual information between the true clustering and the clustering obtained with the trimming procedure with parameter $q = 103$. The confidence intervals for the normalised mutual information are shown in Figure 3.3. Note that in general, the divergences associated to the distributions provide a better clustering, although it is less marked than for data without noise, see for instance the Binomial mixture case.

Generative Model	$d_{Gaussian}$	$d_{Poisson}$	$d_{Binomial}$	d_{Gamma}
Gaussian	0.6585 ± 0.0130	0.6273 ± 0.0113	0.6370 ± 0.0108	0.5926 ± 0.0119
Poisson	0.6582 ± 0.0134	0.6872 ± 0.0126	0.6790 ± 0.0146	0.6500 ± 0.0120
Binomial	0.7295 ± 0.0127	0.7271 ± 0.01381	0.7290 ± 0.0125	0.6990 ± 0.0110
Gamma	0.7653 ± 0.0122	0.8122 ± 0.0110	0.7912 ± 0.0118	0.8142 ± 0.0129
Cauchy	0.7074 ± 0.0149	0.7204 ± 0.0128	0.6876 ± 0.01278	0.6729 ± 0.0116

Figure 3.16: Normalised mutual information – Noisy data

Experiments in higher dimension

We have made the same experiments in dimension d , with the same parameters, by replacing X_i by $(X_i^1, X_i^2, \dots, X_i^d)$ where the X_i^j are all independent and from the same component of the original mixture. In the algorithm, we replace the Bregman divergence with the sum over all coordinates of the 1-dimensional Bregman divergences. For the dimension $d = 2$, we plot the cost in function of the trimming parameter in Figure 3.21 and the corresponding clustering in Figure 3.26. In Figure 3.29 we computed the mutual information for trimming parameter $q = 100$, in Figure 3.30 for trimming parameter $q = 110$ and in Figure 3.31, in dimension 100 for trimming parameter $q = 110$. We proceed like for dimension 1 with 100 points sampled with additional 20 points of noise.




Figure 3.17: Gaussian mixture




Figure 3.18: Poisson mixture




Figure 3.19: Binomial mixture




Figure 3.20: Gamma mixture

Figure 3.21: Trimming parameter q selection - dimension 2




Figure 3.22: Gaussian mixture




Figure 3.23: Poisson mixture




Figure 3.24: Binomial mixture




Figure 3.25: Gamma mixture

Figure 3.26: Clustering associated to the selected parameter q - dimension 2




Figure 3.27: Cauchy mixture




Figure 3.28: Cauchy mixture

Generative Model	$d_{Gaussian}$	$d_{Poisson}$	$d_{Binomial}$	d_{Gamma}
Gaussian	0.6543 ± 0.0107	0.6356 ± 0.0094	0.6426 ± 0.0106	0.5732 ± 0.0017
Poisson	0.7021 ± 0.0098	0.7073 ± 0.0118	0.7107 ± 0.0103	0.6704 ± 0.0089
Binomial	0.7650 ± 0.0116	0.7592 ± 0.0102	0.7728 ± 0.0102	0.7073 ± 0.0083
Gamma	0.7782 ± 0.0082	0.8175 ± 0.0093	0.8041 ± 0.0089	0.7867 ± 0.0079

Figure 3.29: Normalised mutual information – Noisy data - dimension 2 - $q = 100$

Generative Model	$d_{Gaussian}$	$d_{Poisson}$	$d_{Binomial}$	d_{Gamma}
Gaussian	0.6911 ± 0.0087	0.6332 ± 0.0098	0.6537 ± 0.0092	0.5904 ± 0.0141
Poisson	0.7237 ± 0.0095	0.7450 ± 0.0091	0.7432 ± 0.0085	0.7109 ± 0.0085
Binomial	0.7682 ± 0.0073	0.7858 ± 0.0074	0.7784 ± 0.0074	0.7338 ± 0.0089
Gamma	0.7800 ± 0.0073	0.8167 ± 0.0047	0.8038 ± 0.0059	0.8349 ± 0.0043
Cauchy	0.6667 ± 0.0115	0.6375 ± 0.0102	0.6400 ± 0.0119	0.6409 ± 0.0103

Figure 3.30: Normalised mutual information – Noisy data - dimension 2 - $q = 110$

Note that the heuristic for selecting the trimming set is not as effective as for dimension 1. It may come from the fact that the clutter noise is closer to signal in dimension 2. The Bregman divergences for the Poisson, binomial and gamma distributions take into account the

Generative Model	$d_{Gaussian}$	$d_{Poisson}$	$d_{Binomial}$	d_{Gamma}
Gaussian	0.8528 ± 0.0026	0.7868 ± 0.0061	0.7846 ± 0.0065	
Poisson	0.8525 ± 0.0025	0.8690 ± 0.0033	0.8564 ± 0.0038	0.8738 ± 0.0030
Binomial	0.8517 ± 0.0024	0.8694 ± 0.0031	0.8572 ± 0.0038	0.8773 ± 0.0025
Gamma	0.8576 ± 0.0031	0.8727 ± 0.0029	0.8616 ± 0.0033	0.8812 ± 0.0017

Figure 3.31: Normalised mutual information – Noisy data - dimension 100 - $q = 110$

heteroscedasticity of the data, since the variance within a direction (0, 1) or (1, 0) depends on the mean within the direction. In the multidimensional Gaussian setting, Garcia et al. proposed in [GEGMMI08] an algorithm, implemented in the R package tclust, that adapts to mixtures with different covariance matrices. It may be future work to adapt our methods to non-Gaussian distributions with non-diagonal covariance matrices.

3.2 Heteroscedastic Gaussian mixture clustering with the trimmed log-likelihood

In Section 3.1, we exposed a method of clustering for data sampled on \mathbb{R}^d . The method was particularly adapted to mixtures of distributions from some family $\mathcal{F} = \{P_\theta \mid \theta \in \Theta\}$, where P_θ is a probability measure on $\mathcal{X} \subset \mathbb{R}^d$ with density $p_\theta : x \mapsto \exp(\langle x, \theta \rangle - \psi(\theta))$ for every $x \in \mathcal{X}$.

We considered families \mathcal{F} for which the parameter θ was determined by the expectation $P_\theta u$. Such a property is satisfied for all families of Gaussian distributions with a fixed covariance matrix for instance. The clustering method of Section 3.1 is particularly adapted to recover clusters from data sampled according to a mixture of k distributions in \mathcal{F} . However, the method fails drastically when the covariances of the distributions in the mixture are different, as enhanced by Figure 3.35. In this figure, we applied the method presented in Section 3.1.5 to a 250-sample from a mixture of 3 Gaussian distributions with different covariance matrices with additional 50 points of clutter noise. The method does not allow to recover relevant clusters.

In order to overcome this issue, we decide to introduce an additional parameter χ to allow different covariance matrices in the multidimensional Gaussian framework for instance.




Figure 3.32: True labels




Figure 3.33: Trimmed k-means




Figure 3.34: Our new method

Figure 3.35: Trimmed Bregman clustering – Data sampled from a heteroscedastic Gaussian mixture

Moreover, in the previous section we worked with a generalisation of the DTM, the Bregman divergence-to-measure, that expresses from a Bregman divergence d_ϕ as:

$$\inf_{\tilde{P}} \tilde{P} d_\phi(u, \theta),$$

where the infimum is taken over distributions \tilde{P} that are sub-measures of P with P -mass h . The clustering methods introduced were well suited for mixtures of distributions with density of type $p_\theta(x) = \exp(-d_\phi(x, \theta) + C(x))$ for some function C ¹. It is possible to adapt the concept of Bregman divergence-to-measure by replacing $-d_\phi(x, \theta)$ with $-\log(p_\theta(x))$, for some density function p_θ in a family $\mathcal{F} = \{p_\theta \mid \theta \in \Theta\}$.

By doing so, over a first phase, we define the notion of trimmed log-likelihood-to-measure as a continuous version of the trimmed log-likelihood, an alternative to the Bregman divergence-to-measure, or as well, as generalisation of the distance-to-measure. Over a second phase, we introduce a new clustering method based on the trimmed log-likelihood-to-measure to cluster data from mixtures of Gaussian distributions with different covariance matrices, possibly corrupted with clutter noise. This method appears as an alternative to the method proposed in [GEGMMI08]. The algorithm we propose is based on an approach that is studied in more details in the following chapter.

3.2.1 The trimmed log-likelihood - a notion related to the distance to measure

In the following section, we will work with the setting about to be mentioned. The set Θ will be the space of parameters and $(\mathcal{X}, \mathcal{A}, P^0)$ a Polish measured space. To all $\theta \in \Theta$, we will associate P_θ , a probability measure absolutely continuous with respect to P^0 , and denote by p_θ its Radon-Nikodym derivative with respect to P^0 . Also, we will work with a model of probability distributions $\mathcal{F} = \{P_\theta \mid \theta \in \Theta\}$ and a probability distribution P absolutely continuous with respect to P^0 .

The notion of likelihood was first introduced by Fisher in 1922 [Fis22, Section 6] in the context of model estimation. Given $\mathbb{X} = (X_i)_{i \in [1, n]}$, n independent realisations from a distribution P , the *likelihood* is defined for all $\theta \in \Theta$ by $l(\theta) = \prod_{i=1}^n p_\theta(X_i)$. If P is in \mathcal{F} or is a slight modification of a distribution $P_\theta \in \mathcal{F}$, it is of interest to infer θ from the data \mathbb{X} . An estimator of the parameter θ from the likelihood is given by $\hat{\theta}_{MLE} = \arg \max_{\theta \in \Theta} l(\theta)$, and is called *maximum likelihood estimator*. As well, $\hat{\theta}_{MLE}$ is the maximiser of the *log-likelihood* defined by $L(\theta) = \sum_{i=1}^n \log(p_\theta(X_i))$.

The presence of outliers in \mathbb{X} may change drastically the estimator $\hat{\theta}_{MLE}$. Thus, in 1990 Neykov and Neytchev introduced the notion of *trimmed likelihood estimator* in [NN90], as a maximiser of the trimmed likelihood or equivalently to the trimmed log-likelihood, where the *trimmed log-likelihood* is defined for some $k \in [1, n]$ by

$$L(\theta) = \frac{1}{k} \sum_{i=1}^k \log \left(p_\theta \left(X^{(i)} \right) \right), \tag{3.2}$$

where we renumbered the X_i s by $(X^{(i)})_{i \in [1, n]}$ such that:

$$\forall i \leq j, \log \left(p \left(X^{(i)} \right) \right) \geq \log \left(p \left(X^{(j)} \right) \right).$$

Morally, a set of N data points made of k points sampled from a slight modification of a distribution P_{θ_0} in the model \mathcal{F} and $N - k$ outliers will more likely come from a distribution P_θ for $\theta \in \Theta$ when the maximum of the log-likelihood computed from k points is maximised over all subsets of k points among the N data points, and this maximum is attained at θ .

The trimmed likelihood has been extensively investigated, see for instance [HL97], [BC93] and [Mül01] just to name a few. The trimmed log-likelihood estimators are just a particular case

¹This density rewrites as $\exp(\langle x, \theta \rangle - \psi(\theta))$ for some function ψ

of M -estimators introduced in [Hub64] by Huber in 1964. Also, [HL97] in 1997 and [VN98] in 1998 introduced the notion of weighted trimmed log-likelihood, a generalisation of the trimmed log-likelihood with additional weights, that was studied for instance in [Či08], [Mar00] and [NFDN07]. In this work, we focus on the trimmed log-likelihood and introduce a continuous version of it, as follows.

Definition of the trimmed loglikelihood-to-measure

The notion of trimmed log-likelihood can be adapted to general distributions in the following sense.

Definition 3.23. Given a mass parameter $h \in (0, 1]$, the *trimmed log-likelihood-to-measure* P (TLM) is a function defined for all $\theta \in \Theta$ by:

$$L_{P,h}(\theta) = \sup \left\{ \tilde{P} \log(p_\theta(u)) \mid \tilde{P} \in \mathcal{P}_h(P) \right\}. \quad (3.3)$$

The trimmed log-likelihood-to-measure P rewrites as the P -integral of $\log(p_\theta)$ on the upper-level set of P -mass h of the density p_θ , as follows.

Proposition 3.24. For $\theta \in \Theta$ and $h \in (0, 1]$, let $\delta_{P,h}^{\mathcal{F}}(\theta)$ denote the h -upper quantile of the distribution of $\log(p_\theta(X))$ when X is sampled according to P , that is:

$$\delta_{P,h}^{\mathcal{F}}(\theta) = \sup \left\{ r \in \mathbb{R} \mid P \left(\overline{\mathcal{B}}^{\mathcal{F}}(\theta, r) \right) > h \right\}.$$

where $\mathcal{B}^{\mathcal{F}}(\theta, r) = \{x \in \mathcal{X} \mid \log(p_\theta(x)) > r\}$ stands for the r -super-level set of the logarithm of the density p_θ for $r \in \mathbb{R}$, and $\overline{\mathcal{B}}^{\mathcal{F}}(\theta, r)$ for its closure. Note that $\overline{\mathcal{B}}^{\mathcal{F}}(\theta, -\infty)$ coincides with the whole space \mathcal{X} .

Then, by denoting

$$\mathcal{P}_{\theta,h}(P) = \left\{ \tilde{P} \in \mathcal{P}_h(P) \mid h\tilde{P} = P \text{ on } \mathcal{B}^{\mathcal{F}}(\theta, \delta_{P,h}^{\mathcal{F}}(\theta)) \text{ and } \text{Supp}(\tilde{P}) \subset \overline{\mathcal{B}}^{\mathcal{F}}(\theta, \delta_{P,h}^{\mathcal{F}}(\theta)) \right\}, \quad (3.4)$$

for every $P_{\theta,h} \in \mathcal{P}_{\theta,h}(P)$, we have that:

$$L_{P,h}(\theta) = P_{\theta,h} \log(p_\theta(u)). \quad (3.5)$$

Proof. The following Lemma is inspired by [CCSM11, Proposition 3.3]. It is a direct adaptation of the proof of Lemma 3.5 with $F_\theta^{-1}(u)$ the u -upper quantile of $\log(p_\theta(X))$ for $u \in [0, 1]$ and X , a random variable from P . \square

So, the trimmed log-likelihood-to-measure, associated to a distribution P , at a point $\theta \in \Theta$, is denoted by $L_{P,h}(\theta)$. It is defined as the supremum (among all sub-measures of P with P -mass h) of the expectation of the log-likelihood $\log(p_\theta)$ according to the sub-measure in question. This supremum is attained at the sub-measure coinciding with P on the sub-level set of p_θ with P -mass h . When \mathcal{F} is the Gaussian family, this sub-level set is a ball. Then, the trimmed log-likelihood-to-measure and the DTM coincide (up to some constants).

When X_1, X_2, \dots, X_n is an n -sample from P , and P_n denotes the empirical measure associated with the X_i s, the TLM $L_{P_n,h}$ coincides with the trimmed log-likelihood. Indeed, when $h = \frac{k}{n}$ with $k \in \llbracket 1, n \rrbracket$:

$$L_{P_n,h}(\theta) = \frac{1}{k} \sum_{i=1}^k \log \left(p_\theta \left(X^{(i)} \right) \right). \quad (3.6)$$

Expression of the trimmed loglikelihood-to-measure in terms of the distance-to-measure

As noted in [HL97], [BC93] or [Mül01] just to name a few, the notion of trimmed log-likelihood estimator (that they call location estimator) is to relate with the notion of least trimmed squares estimators as defined in [Rou84], [Rou85] and [RL87] when the family of distributions considered is the family of Gaussian distributions with fixed covariance matrix. The analogous consideration for general distributions P is that the notion of trimmed log-likelihood-to-measure is to relate with the notion of distance-to-measure as defined in [CCSM11] also known as the h -trimmed variation of [CAGM97], also see Definition 3.4. Indeed, for a family \mathcal{F} of probabilities with density $p_\theta(x) = C \exp(-d^2(x, \theta))$ for some metric d on \mathcal{X} and normalisation constant $C > 0$, the log-likelihood rewrites in function of the distance to the measure P for the metric d (denoted by $d_{P,h,d}$) as follows:

$$L_{P,h}(\theta) = \log(C) - d_{P,h,d}^2(\theta). \quad (3.7)$$

Indeed, the upper-level sets of the density p_θ are balls centred at θ for the metric d , thus, the set of sub-measures $\mathcal{P}_{\theta,h}(P)$ defined by (3.4) coincides with the set of sub-measures for the definition of the distance-to-measure with the metric d , see Definition 2.1.

The example of the family of multivariate Gaussian distributions

As an example, when \mathcal{F} denotes the family of multivariate Gaussian distributions with a fixed covariance matrix Σ , Equation (3.7) is satisfied for $C = \frac{1}{\sqrt{(2\pi)^d \det(\Sigma)}}$ and $d(x, \theta) = \frac{1}{2} \|x - \theta\|_{\Sigma^{-1}}$, where $\|\cdot\|_{\Sigma^{-1}}$ denotes the Σ -Mahalanobis metric and is defined for all x and θ in \mathbb{R}^d by $\|x - \theta\|_{\Sigma^{-1}} = \sqrt{(x - \theta)^T \Sigma^{-1} (x - \theta)}$. When $\Sigma = I_d$ is the identity matrix in \mathbb{R}^d , the Mahalanobis metric coincides with the Euclidean metric. Moreover, the trimmed log-likelihood associated with an *isotropic Gaussian model* $\mathcal{F} = \{P_{(\mu,\sigma)} = \mathcal{N}(\mu, \sigma^2 I_d) \mid \mu \in \mathbb{R}^d, \sigma > 0\}$ is proportional to the squared distance-to-measure associated with the Euclidean distance, $d_{P,h,\|\cdot\|}^2$, as defined in [CCSM11]. Indeed, we have that

$$L_{P,h}((\mu, \sigma)) = -\frac{d}{2} \log(2\pi) - \frac{d}{2} \log(\sigma^2) - \frac{1}{2\sigma^2} d_{P,h,\|\cdot\|}^2(\mu). \quad (3.8)$$

It turns out that, in the Gaussian context, the mean and the covariance matrix of any distribution P_{θ_0} in the model are recovered from the trimmed likelihood-to-measure maximiser, up to a normalisation with the distance to the measure $\mathcal{N}(0, I_d)$. We start with the isotropic case.

Proposition 3.25. *If $\mathcal{F} = \{\mathcal{N}(\mu, \sigma^2 I_d) \mid \mu \in \mathbb{R}^d, \sigma > 0\}$ and $P = \mathcal{N}(\mu_0, \sigma_0^2 I_d)$, then $L_{P,h}((\mu, \sigma))$ is maximal for $\mu = \mu_0$ and $\sigma^2 = \sigma_0^2 \frac{d_{\mathcal{N}(0, I_d), h, \|\cdot\|}^2(0)}{d}$.*

Proof. When \mathcal{F} is a family of isotropic Gaussian distributions, it comes from (3.8) that the maximum of the trimmed log-likelihood to the measure P is attained at $\sigma^2 = \frac{d_{P,h,\|\cdot\|}^2(\mu)}{d}$ and $\mu \in \mathbb{R}^d$, the minimiser of $\mu \mapsto d_{P,h,\|\cdot\|}(\mu)$. Note that whenever $P = \mathcal{N}(\mu_0, \sigma_0^2 I_d)$ for some $\sigma_0 > 0$, this minimum is attained at $\mu = \mu_0$ and $\sigma^2 = \sigma_0^2 \frac{d_{\mathcal{N}(0, I_d), h, \|\cdot\|}^2(0)}{d}$. Indeed, $d_{P,h,\|\cdot\|}(\mu)$ rewrites $\sigma_0 d_{\mathcal{N}(0, I_d), h, \|\cdot\|}(\mu - \mu_0)$, since $x \in \mathcal{B}(\mu_0, \delta_{P,h}(\mu_0))$ if and only if $\frac{x - \mu_0}{\sigma_0} \in \mathcal{B}(0, \delta_{\mathcal{N}(0, I_d), h}(0))$. \square

When \mathcal{F} is a family of multivariate Gaussian distributions $P_{\mu,\Sigma}$ with covariance matrix $\Sigma = AA^T$, the trimmed log-likelihood to P at (μ, Σ) rewrites

$$L_{P,h}((\mu, \Sigma)) = -\frac{d}{2} \log(2\pi) - \frac{1}{2} \log(\det(\Sigma)) - \frac{1}{2} d_{P,h,\|\cdot\|_{\Sigma^{-1}}}^2(\theta). \quad (3.9)$$

The analogous result of Proposition 3.25 for the non isotropic case is the following.

Proposition 3.26. *If $\mathcal{F} = \{\mathcal{N}(\mu, \Sigma) \mid \mu \in \mathbb{R}^d, \Sigma \text{ symmetric definite positive}\}$ and $P = \mathcal{N}(\mu_0, \Sigma_0)$, then $L_{P,h}((\mu, \Sigma))$ is maximal for $\mu = \mu_0$ and $\Sigma = \Sigma_0 \frac{d_{\mathcal{N}(0, I_d), h, \|\cdot\|}^{(0)}}{d}$.*

The proof of Proposition 3.26 is deferred to Section 3.3.2.

It is based on the following Lemma. I am grateful to Luc Lehericy for his help.

Lemma 3.27. *For all $h \in [0, 1]$, the distance to the measure $\mathcal{N}(0, I_d)$ at 0 for the metric $\|\cdot\|_{\Sigma^{-1}}$, $d_{\mathcal{N}(0,1), h, \|\cdot\|_{\Sigma^{-1}}}$ is minimal for $\Sigma = I_d$ among all symmetric positive definite matrices Σ with determinant 1. The matrix I_d is the only minimizer when $h \in (0, 1)$.*

The proof of Lemma 3.27 is to be found in Section 3.3.2.

Rates of convergence for the trimmed log-likelihood-to-measure

As aforementioned, the notion of trimmed log-likelihood is strongly related to the notion of distance-to-measure. Some properties for the trimmed log-likelihood can be derived from the corresponding properties satisfied by the distance-to-measure. In particular, the trimmed log-likelihood associated with a probability measure P , $L_{P,h}$ can be approximated by its empirical counterpart, $L_{P_n, h}$. And it is possible to obtain deviation bounds for $L(\mathcal{K}, P_n) = \sup_{\theta \in \mathcal{K}} \|L_{P,h}(\theta) - L_{P_n, h}(\theta)\|$ for \mathcal{K} a compact subset of Ω under some assumptions on \mathcal{F} and P .

For instance, when \mathcal{K} is reduced to a single element, it is possible to derive rates of convergence for $\mathbb{E}[L(\mathcal{K}, P_n)]$ in $\frac{1}{\sqrt{n}}$ or even faster rates when h decreases to zero when n goes to infinity and also exponential deviation bounds for $L(\mathcal{K}, P_n)$. The rates will depend on the continuity modulus of $\theta \mapsto \delta_{P,h}^{\mathcal{F}}(\theta)$ if we adapt the proofs of [CMM16]. In order to derive rates of convergence for compact sets \mathcal{K} not necessarily reduced to a single element, the bound will depend on the covering number of $\theta \mapsto L_{P,h}(\theta)$.

When p_θ rewrites like $x \mapsto \exp(\langle \theta, f(x) \rangle - \psi(\theta))$, with $\psi \in \mathcal{C}_1$ and finite on the support of the measure P , $\tilde{M}_1 = Pf(u) < +\infty$ and $M_\infty = \sup_{\theta \in \mathcal{K}} \sup_{x \in \text{Supp}(P)} \log(p_\theta(x)) < +\infty$ then $\sup_{\mathcal{K}} |L_{P_n, h}(\theta) - L_{P, h}(\theta)|$ is of order $\frac{1}{\sqrt{n}}$.

Proof. Roughly, the standard rate can be directly derived from the method in Section 3.1. Indeed, the proof of Lemma 3.59 adapts directly since the balls $\mathcal{B}^{\mathcal{F}}(\theta, r)$ are actually half-spaces, thus with Vapnik-Chervonenkis dimension at most $d' + 1$ when $\Theta \subset \mathbb{R}^{d'}$. It yields that

$$\sup_{\theta \in \mathcal{K}, r \in \mathbb{R}} \left| (P_n - P) \mathbb{1}_{\mathcal{B}^{\mathcal{F}}(\theta, r)}(u) \right| \leq C \frac{\sqrt{d' + 1}}{\sqrt{n}}$$

with probability larger than $\exp(-p)$. Also, Lemma 3.60 adapts and entails that with probability larger than $\exp(-p)$,

$$\sup_{\theta \in \mathcal{K}, r \in [r_0, +\infty)} \left| (P_n - P) \log(p_\theta(u)) \mathbb{1}_{\mathcal{B}^{\mathcal{F}}(\theta, r)}(u) \right| \leq C_{\mathcal{K}, \tilde{M}_1} \max(r_0, M_\infty) \frac{\sqrt{d' + 1}}{\sqrt{n}} + \max(r_0, M_\infty) \sqrt{\frac{2p}{n}}.$$

Indeed, $|\log(p_\theta(x)) - \log(p_{\theta'}(x))| \leq (\|f(x)\| + \sup_{\theta \in \mathcal{K}} |\nabla_\theta \psi|) \|\theta - \theta'\|$ according to the mean value theorem. Thus,

$$\begin{aligned} \sup_{\mathcal{K}} |L_{P_n, h}(\theta) - L_{P, h}(\theta)| &= |P \log(p_\theta(u)) \mathbb{1}_{\mathcal{B}(\theta, \delta_{P, h}^{\mathcal{F}}(\theta))}(u) - P_n \log(p_\theta(u)) \mathbb{1}_{\mathcal{B}(\theta, \delta_{P_n, h}^{\mathcal{F}}(\theta))}(u)| \\ &\leq M_\infty P_n |\mathbb{1}_{\mathcal{B}(\theta, \delta_{P_n, h}^{\mathcal{F}}(\theta))}(u) - \mathbb{1}_{\mathcal{B}(\theta, \delta_{P, h}^{\mathcal{F}}(\theta))}(u)| + |(P - P_n) \log(p_\theta(u)) \mathbb{1}_{\mathcal{B}(\theta, \delta_{P, h}^{\mathcal{F}}(\theta))}(u)| \\ &\leq M_\infty |(P_n - P) \mathbb{1}_{\mathcal{B}(\theta, \delta_{P, h}^{\mathcal{F}}(\theta))}(u)| + |(P - P_n) \log(p_\theta(u)) \mathbb{1}_{\mathcal{B}(\theta, \delta_{P, h}^{\mathcal{F}}(\theta))}(u)| \\ &\leq C_{\mathcal{K}, \tilde{M}_1} \max(r_0, M_\infty) \frac{\sqrt{d' + 1}}{\sqrt{n}} + \max(r_0, M_\infty) \sqrt{\frac{2p}{n}} \end{aligned}$$

with probability larger than $2 \exp(-p)$ since $P_n(\mathcal{B}(\theta, \delta_{P_n, h}^{\mathcal{F}}(\theta))) = h = P(\mathcal{B}(\theta, \delta_{P, h}^{\mathcal{F}}(\theta)))$, with $r_0 = \inf_{\theta \in \mathcal{K}} \delta_{P, h}^{\mathcal{F}}(\theta) < -\infty$. \square

It may also be possible to adapt the proof of Theorem 3.15 to derive deviation bounds for $|L_{P, h}(\hat{\theta}_n) - L_{P, h}(\theta^*)|$ where $\hat{\theta}_n$ is the minimiser of $L_{P_n, h}$ and θ^* the minimiser of $L_{P, h}$.

Such results complete the numerous work to assess the quality of the trimmed likelihood estimators. For instance, in [?], the authors prove that maximum likelihood estimates based on any trimming of the sample space suffice to successfully estimate the parameters of a Gaussian distribution. A review on asymptotic efficiency of log-likelihood maximum estimators is to be found in [AT11]. In [VDG96], van de Geer derives rates of convergence for the log-likelihood estimator in mixture models in the sense of the Hellinger distance to the true density. Also see [JJN85].

According to [CAGM97, Theorem 3.6], for families that rewrite as $p_\theta(x) = \exp(-\phi(\|x - \theta\|))$ for $\theta \in \mathcal{X}$ and for some convex non-decreasing function ϕ , if P is continuous and the maximiser θ^* of $L_{P, h}(\theta)$ is unique, then the maximiser $\hat{\theta}$ of $L_{P_n, h}(\theta)$ converges almost surely to θ^* .

The theorem enlarges to mixtures of distributions of this type for $\theta^* = (\theta_1^*, \theta_2^*, \dots, \theta_k^*)$ a maximiser of a slightly modified version of the trimmed log-likelihood to the measure P . Since then, the maximisation of the log-likelihood method boils down to a trimmed k -means as defined in [CAGM97].

In the following, we propose a generalisation of the trimmed clustering for data sampled from a mixture of k distributions with different covariance matrices, corrupted by noise.

3.2.2 A new clustering method based on the trimmed log-likelihood and a modified version of the Lloyd's algorithm

Given a model \mathcal{F} , a mixture of k distributions in \mathcal{F} is a probability P defined by $P = \sum_{i=1}^k \alpha_i P_{\theta_i}$ for some non-negative $\alpha_1, \alpha_2, \dots, \alpha_k$ with $\sum_{i=1}^k \alpha_i = 1$. The corresponding mixture model is defined by $\mathcal{F}^k = \left\{ \sum_{i=1}^k \alpha_i P_{\theta_i} \mid \alpha_i \geq 0, \sum_{i=1}^k \alpha_i = 1 \right\}$.

A reference document on mixture models is the book from McLachlan and Deep, see [MP00]. A reference on mixture of exponential families, maximum likelihood and the EM algorithm is given by [RW84]. The particular case of multidimensional Gaussian mixture models with the maximum log-likelihood and the EM algorithm is well understood, see for instance [CG95]. The case of data corrupted with clutter noise is treated in [DR98] where for instance, features are detected. See [GEGMM10] for a review on trimmed clustering. I am grateful to Giles Celeux for the references.

The literature relative to Gaussian mixture recovering is proficient, see [BCRR97], [CM74], [GR05], [FR98], [Hen04], [PM00] and [Ste00].

The trimmed log-likelihood has been used on the context of mixture models recovering, see [Mar00], and [NFDN07] with the many references therein. In [CAMMI08] Cuesta et al. consider the special case of Gaussian mixtures. This is still an active research domain, see for instance [YYC16] where the authors proposed a new version of the EM algorithm based on the trimmed likelihood method and M estimation approaches. See also [TK18].

The method we propose is an alternative to the work in [GEGMMI08] also based on the trimmed log-likelihood, with asymptotic results, or [DFGEMI18]. Note that most of the methods proposed for trimmed clustering under the heteroscedastic assumption are based on the

EM algorithm, the method we propose in the following is not, it is in fact a generalisation of Algorithm 3.1 based on the trimmed k-means method. Another generalisation of Algorithm 3.1 will be used in the next chapter to another purpose.

In this section, we will work with the following setting. We set \mathcal{X} , a subset of Euclidean space \mathbb{R}^d equipped with a measure P^0 , Θ a convex subset of $\mathbb{R}^{d'}$, defined by

$$\Theta = \{\theta(\mu, \chi) \mid \mu \in \mathcal{X}, \chi \in \Xi\}.$$

The parameter space Θ is itself parametrised by \mathcal{X} and some parameter space Ξ , included in some Euclidean space. Concretely, we parametrise θ with an element μ of \mathcal{X} that will play the role of the mean $P_\theta u$, and an additional parameter $\chi \in \Xi$ for taking into account the possible heteroscedasticity of the distributions in the mixture model. We assume that $\mu \mapsto \Theta(\mu, \chi)$ is linear, so that $P\theta(u, \chi) = \theta(Pu, \chi)$.

We consider an exponential family $\mathcal{F}_{f,\psi} = \{P_\theta \mid \theta \in \Theta\}$ parametrized by a sufficient statistic f that is a measurable function on \mathcal{X} taking values in $\mathbb{R}^{d'}$, and a proper closed strictly convex cumulant function ψ defined on Θ as follows. For all $\theta \in \Theta$, P_θ is absolutely continuous with respect to P^0 , with Radon-Nikodym density given by:

$$p_\theta(x) = \exp(\langle \theta, f(x) \rangle - \psi(\theta)). \quad (3.10)$$

As aforementioned, the Bregman divergence-to-measure also known as h -trimmed variation, see Section 3.1.3, is particularly adapted for clustering data sampled from a mixture of distributions in $\mathcal{F}_{f,\psi}$ when $\theta = \theta(\mu)$ is parametrised by the expectation $\mu = P_\theta u$. Indeed, for such distributions, when $f(x) = x$, the density rewrites as $p_\theta(x) = \exp(-d_\phi(x, \theta) + \phi(x))$ for some convex function ϕ .

The remaining clustering consists in a union of k Bregman-balls with the same radius, that is, an upper-level set of the density of a mixture of k distributions in $\mathcal{F}_{f,\psi}$. As enhanced by Figure 3.35, the method fails for a mixture of Gaussian with different covariance matrices.

We might have the idea to adapt the method by replacing the $d_\phi(x, \theta)$ with $-\log(p_\theta(x))$ in the definition of the Bregman divergence-to-measure at a centre θ . The reader should be aware of the fact that Algorithm 3.1 cannot be directly adapted, indeed, $(x, \mu) \mapsto -\log(p_{\theta(\mu)}(x))$ is no longer a Bregman divergence. In particular, the property that $P - \log(p_\theta(u))$ is minimal at $\theta = \theta(Pu)$ is no longer satisfied, since such a property is only satisfied for Bregman divergences, according to [BGW05]. Thus, the step in the algorithm consisting in replacing θ_i with the mean of the sample points within the i -th Bregman Voronoï cell will not necessarily decrease the loss $P_n \min_{i \in \llbracket 1, k \rrbracket} -\log(p_{\theta_i}(u))$.

We will adapt Algorithm 3.1 with a well-chosen Bregman divergence depending on covariance matrices $(\Sigma_i)_{i \in \llbracket 1, k \rrbracket}$ that are updated at each step of the algorithm.

The method we propose adapts to not well-separated clusters, we will see that it improves the method of [GEGMMI08] implemented with the R package tclust.

Some theory

For $P \in \mathcal{P}(\mathcal{X})$, $\theta \in \Theta$, $h \in (0, 1)$ and the model $\mathcal{F}_{f,\psi} = \{P_\theta \mid \theta \in \Theta\}$ of probabilities with density defined by (3.10), the set $\mathcal{P}_{\theta,h}(P)$ defined by (3.4) can be rewritten as

$$\mathcal{P}_{\theta,h}(P) = \arg \max \left\{ \langle \theta, \tilde{P} f(u) \rangle \mid \tilde{P} \in \mathcal{P}_h(P) \right\}. \quad (3.11)$$

We still denote by $P_{\theta,h}$ any element in $\mathcal{P}_{\theta,h}(P)$.

Concretely, by denoting

$$c_{P,h}(\theta) = \sup \{c \in \mathbb{R} \mid P(H(\theta, c)) > h\},$$

with

$$H(\theta, c) = \{x \in \mathcal{X} \mid \langle f(x), \theta \rangle > c\},$$

and $\bar{H}(\theta, c)$ its closure, $\mathcal{P}_{\theta,h}(P)$ exactly corresponds to the set of all distributions $P_{\theta,h} = \frac{1}{h}\tilde{Q} \in \mathcal{P}_h(P)$ supported on $\bar{H}(\theta, c_{P,h}(\theta))$ and such that $\tilde{Q} = P$ on $H(\theta, c_{P,h}(\theta))$.

According to Proposition 3.24, the TLM satisfies for all $\theta \in \Theta$:

$$L_{P,h}(\theta) = \sup_{\tilde{P} \in \mathcal{P}_h(P)} \langle \theta, \tilde{P}f(u) \rangle - \psi(\theta) \quad (3.12)$$

$$= \sup_{\theta' \in \Theta} \langle \theta, P_{\theta',h}f(u) \rangle - \psi(\theta) \quad (3.13)$$

$$= \langle \theta, P_{\theta,h}f(u) \rangle - \psi(\theta). \quad (3.14)$$

In the following we propose to derive an approximation of $L_{P,h}$ given by:

$$L_{P,h,k}(\theta) = \max_{i \in \llbracket 1, k \rrbracket} \langle \theta, P_{\theta_i^*, h}f(u) \rangle - \psi(\theta),$$

with

$$\theta^* \in \arg \max \left\{ Pf(\theta, u) \mid \theta = (\theta_i = \theta(\mu_i, \chi_i))_{i \in \llbracket 1, k \rrbracket} \in \Theta^{(k)} \right\}, \quad (3.15)$$

with

$$f(\theta, x) = \max_{i \in \llbracket 1, k \rrbracket} \langle \theta(x, \chi_i), P_{\theta_i, h}f(u) \rangle - \psi(\theta(x, \chi_i)).$$

Proposition 3.28. *Let $P = P_{\theta_0}$ for some $\theta_0 \in \Theta$, then for $k = 1$ the optimum θ^* satisfies:*

$$\theta^* = P\theta(u, \chi^*) = \theta(Pu, \chi^*)$$

with

$$\chi^* \in \arg \max \left\{ \langle P\theta(u, \chi), P_{P\theta(u, \chi), h}f(u) \rangle - P\psi(\theta(u, \chi)) \right\}.$$

Proof. For all $\theta = \theta(\mu, \chi) \in \Theta$, by definition of the measures $P_{\theta,h}$, according to (3.13) and (3.14), we have:

$$\begin{aligned} Pf(\theta, u) &= \langle P\theta(u, \chi), P_{\theta,h}f(u) \rangle - P\psi(\theta(u, \chi)) \\ &\leq \langle P\theta(u, \chi), P_{P\theta(u, \chi), h}f(u) \rangle - P\psi(\theta(u, \chi)). \end{aligned}$$

□

Note that for all $k \in \mathbb{N}^*$, the maximisation problem given by (3.16) boils down to the following maximisation problem

$$\{\theta^*, \chi^*\} \in \arg \max \left\{ \sup_{\chi = (\chi_i)_{i \in \llbracket 1, k \rrbracket} \in \Xi^{(k)}} Pf(\theta, \chi, u) \mid \theta \in \Theta^{(k)}, \chi \in \Xi^{(k)} \right\}, \quad (3.16)$$

with

$$f(\theta, \chi, x) = \max_{i \in \llbracket 1, k \rrbracket} \langle \theta(x, \chi_i), P_{\theta_i, h}f(u) \rangle - \psi(\theta(x, \chi_i)).$$

Indeed, we may write

$$\begin{aligned}
& P \max_{i \in [1, k]} \langle \theta(v, \chi_i), P_{\theta_i, h} f(u) \rangle - \psi(\theta(v, \chi_i)) dv \\
&= \sum_{i=1}^k P_i(\mathcal{X}) \left\langle \frac{P_i}{P_i(\mathcal{X})} \theta(v, \chi_i) dv, P_{\theta_i, h} f(u) \right\rangle - P_i \psi(\theta(v, \chi_i)) dv \\
&\leq P \max_{i \in [1, k]} \langle \theta(v, \chi_i), P_{\theta'_i, h} f(u) \rangle - \psi(\theta(v, \chi_i)) dv
\end{aligned}$$

where $\theta'_i = \frac{P_i}{P_i(\mathcal{X})} \theta(u, \chi_i) = \theta\left(\frac{P_i}{P_i(\mathcal{X})} u, \chi_i\right)$, and $P = \sum_{i=1}^k P_i$ with the support of P_i included in the set of points $\mu \in \mathcal{X}$ such that $\langle \theta(\mu, \chi_i), P_{\theta_i, h} f(u) \rangle - \psi(\theta(\mu, \chi_i)) \geq \langle \theta(\mu, \chi_j), P_{\theta_j, h} f(u) \rangle - \psi(\theta(\mu, \chi_j))$ for all $j \neq i$. Thus, for all θ_i , there is a better θ'_i which associated χ is equal to χ_i .

Intuitively, θ^* is associated with the best approximation of $L_{P, h}$ from below with the upper envelop of k linear functions $L_i : \theta \rightarrow \langle \theta, f_i^* \rangle - \psi(\theta)$ with f_i^* that maximises the expectation of $\max_{i \in [1, k]} \langle \theta(X, \chi_i^*), f_i \rangle - \psi(\theta(X, \chi_i^*))$ for $X \sim P$. The density of Gaussian distributions rewrite as (3.10) and the maximiser θ^* of the criterion (3.16) can be computed when P is in the model. Indeed, we recover the mean and the covariance matrix up to a constant depending on the distance to the measure $\mathcal{N}(0, I_d)$.

Proposition 3.29. *When $\mathcal{F} = \{\mathcal{N}(\mu, \sigma^2) \mid \mu \in \mathbb{R}, \sigma > 0\}$ and $k = 1$, if $P = \mathcal{N}(\mu_0, \sigma_0^2)$, the optimal parameters are given by $\mu^* = \mu_0$ and $\chi^* = \sigma^{*2} = \left(1 + d_{\mathcal{N}(0,1), h}^2(0)\right) \sigma_0^2$.*

The proof of Proposition 3.29 is to be found in Section 3.3.2.

Thus, the optimal (μ^*, χ^*) determines the distribution $P_{\theta(\mu_0, \chi_0)}$ when dealing with univariate Gaussian distributions. This result enlarges to the dimension d for general covariance matrices.

Proposition 3.30. *When $\mathcal{F} = \{\mathcal{N}(\mu, \Sigma) \mid \mu \in \mathbb{R}^d, \Sigma \text{ symmetric definite positive}\}$ and $k = 1$, if $P = \mathcal{N}(\mu_0, \Sigma_0)$, the optimal parameters are given by $\mu^* = \mu$ and $\chi^* = \Sigma^* = \frac{d + d_{\mathcal{N}(0, I_d), h}^2(0)}{d} \Sigma_0$.*

The proof of Proposition 3.30 can be found in Section 3.3.2. It is possible to derive an algorithm that is a modified version of the Lloyd's algorithm, in order to compute a local maximiser for the criterion (3.16).

Some algorithms

We start with an algorithm that can be applied to every family of measures \mathcal{F} satisfying (3.10), provided that the maximizer of (3.28) can be computed.

Algorithm 3.3: Heteroscedastic clustering – General case

```

Input :  $\mathbb{X}_n$  an  $n$ -sample from  $P$ ,  $q$  and  $k$  ;

# Initialization
Sample  $\mu_1, \mu_2, \dots, \mu_k$  from  $\mathbb{X}_n$  without replacement. ;
Choose arbitrarily  $\eta_1, \eta_2, \dots, \eta_k$  in  $\Xi$  ;
Choose arbitrarily  $\chi_1, \chi_2, \dots, \chi_k$  in  $\Xi$  ;
for  $i$  in  $1..k$ :
     $\theta_i = \theta(\mu_i, \eta_i)$  ;

```

```

while the  $(\theta_i, \chi_i)$ s vary make the following two steps:
  # Decomposition in cells.
  for  $i$  in  $1..k$ :
     $\mathcal{C}(\theta_i, \chi_i) = \emptyset$ ;
  for  $j$  in  $1..n$ :
    Add  $X_j$  to the  $\mathcal{C}(\theta_i, \chi_i)$  (for  $i$  as small as possible) satisfying for all  $l \neq i$ :
     $\langle \theta(X_j, \chi_i), P_{\theta_i, h} f(u) \rangle - \psi(\theta(X_j, \chi_i)) \geq \langle \theta(X_j, \chi_l), P_{\theta_l, h} f(u) \rangle - \psi(\theta(X_j, \chi_l))$ ;
  # Computation of the new centres and weights.
  for  $i$  in  $1..k$ :
     $\theta'_i = \left( \frac{1}{|\mathcal{C}(\theta_i, \chi_i)|} \sum_{X \in \mathcal{C}(\theta_i, \chi_i)} \theta(X, \chi_i) \right)$ ;
     $\chi'_i \in \arg \max \left\{ \frac{1}{|\mathcal{C}(\theta_i, \chi_i)|} \sum_{X \in \mathcal{C}(\theta_i, \chi_i)} \langle \theta(X, \chi), P_{\theta'_i, h} f(u) \rangle - \psi(\theta(X, \chi)) \mid \chi \in \mathcal{X} \right\}$ ;
     $\theta_i = \left( \frac{1}{|\mathcal{C}(\theta_i, \chi_i)|} \sum_{X \in \mathcal{C}(\theta_i, \chi_i)} \theta(X, \chi'_i) \right)$ ;
     $\chi_i = \chi'_i$ ;

Output:  $(\theta_i)_{i \in [1, k]}$ ,  $(\chi_i)_{i \in [1, k]}$ ,  $(\mathcal{C}(\theta_i, \chi_i))_{i \in [1, k]}$ 

```

Proposition 3.31. *The algorithm 3.3 converges to a local maximum of*

$$(\boldsymbol{\mu}, \boldsymbol{\chi}) \in (\mathcal{X} \times \Xi)^{(k)} \mapsto P_n \max_{i \in [1, k]} \langle \theta(v, \chi_i), P_{n \theta_i, h} f(u) \rangle - \psi(\theta(v, \chi_i)) dv,$$

where $\theta_i = \theta(\mu_i, \chi_i)$ for $i \in [1, k]$.

The proof of Proposition 3.31 is deferred to Section 3.3.2.

Algorithm 3.3 applies to the Gaussian model in dimension 1.

The distribution $P_{n \theta(\mu, \chi), h}$ does not depend on $\chi \in \Xi$, thus, we use the notation $P_{n \mu, h}$ for $P_{n \theta(\mu, \chi), h}$. When $h = \frac{q}{n}$, $P_{n \mu, h}$ is the distribution of $\frac{1}{q} \sum_{i=1}^q \delta_{X_i(\mu)}$ where $X_i(\mu)$ is one of the i -th nearest neighbour of μ in \mathbb{X}_n , with mean $c(\mu) = \frac{1}{q} \sum_{i=1}^q X_i(\mu)$ and variance $x^2(\mu) = \frac{1}{q} \sum_{i=1}^q (X_i(\mu) - c(\mu))^2$. We use the notation $|\mathcal{C}|$ for the cardinal of \mathcal{C} .

Finally, note that for every $Q \in \mathcal{P}(\mathcal{X})$, $\mu \in \mathcal{X}$ and $\chi \in \Xi$,

$$Q \langle \theta(v, \chi), P_{n \mu, h}(f(u)) \rangle - \psi(\theta(v, \chi)) dv = -\frac{1}{2} \log(2\pi) - \frac{x^2(\mu) + (Qu - c(\mu))^2 + v(Q)}{2\chi} - \frac{1}{2} \log(\chi),$$

and attains its maximum at $\chi = x^2(\mu) + (Qu - c(\mu))^2 + v(Q)$.

Thus, Algorithm 3.3 adapts into Algorithm 3.4 as follows.

Algorithm 3.4: Heteroscedastic clustering for unidimensional Gaussian mixtures

```

Input:  $\mathbb{X}_n$  an  $n$ -sample from  $P$ ,  $q$  and  $k$  ;
# Initialization
Sample  $\mu_1, \mu_2, \dots, \mu_k$  from  $\mathbb{X}_n$  without replacement ;
for  $i$  in  $1..k$ :
   $\sigma_i = 1$ ;
while the  $\mu_i$ s vary make the following two steps:
  # Decomposition in cells.
  for  $i$  in  $1..k$ :
     $\mathcal{C}_i = \emptyset$ ;
  for  $j$  in  $1..n$ :

```

Add X_j to the \mathcal{C}_i (for i as small as possible) satisfying

$$\frac{\sigma_i^2 \log(\sigma_i^2) + (X_j - c(\mu_i))^2 + x^2(\mu_i)}{\sigma_i^2} \leq \frac{\sigma_l^2 \log(\sigma_l^2) + (X_j - c(\mu_l))^2 + x^2(\mu_l)}{\sigma_l^2} \quad \forall l \neq i;$$

Computation of the new centres and weights.

for i **in** $1..k$:

$$\mu'_i = \frac{1}{|\mathcal{C}_i|} \sum_{X \in \mathcal{C}_i} X;$$

$$v_i = \frac{1}{|\mathcal{C}_i|} \sum_{X \in \mathcal{C}_i} (X - \mu'_i)^2;$$

$$\sigma_i^2 = (c(\mu'_i) - \mu'_i)^2 + v_i + x^2(\mu'_i);$$

$$\mu_i = \mu'_i;$$

Output: $(\mu_i)_{i \in [1,k]}$, $(\sigma_i^2)_{i \in [1,k]}$, $(\mathcal{C}_i)_{i \in [1,k]}$

As a direct corollary of Proposition 3.31 it comes that:

Corollary 3.32. Algorithm 3.4 converges to a local maximum of

$$(\boldsymbol{\mu}, \boldsymbol{\sigma}^2) \mapsto P_n \max_{i \in [1,k]} \log \left(p_{n, \mu_i, \sigma_i^2}(u) \right),$$

where for every $x \in \mathbb{R}$,

$$p_{n, \mu_i, \sigma_i^2}(x) = \frac{1}{\sqrt{2\pi\sigma_i^2}} \exp \left(-\frac{(x - m(P_{n, \mu_i, h}))^2 + v(P_{n, \mu_i, h})}{2\sigma_i^2} \right),$$

for $P_{n, \mu_i, h} \in \mathcal{P}_{\mu_i, h}(P_n)$.

Algorithm 3.3 also adapts to the Heteroscedastic multidimensional Gaussian model. Indeed, the family of distributions $\mathcal{N}(\mu, \Sigma)$ satisfies (3.10) with $\theta(\mu, \Sigma) = (\Sigma^{-1}\mu, \text{vect}(-\frac{1}{2}\Sigma^{-1}))^T$, $f(x) = (x, \text{vect}(xx^T))^T$ and $\psi(\theta) = \frac{1}{2}\mu^T \Sigma^{-1}\mu + \frac{1}{2} \log(\det(2\pi\Sigma))$.

For every $\theta = \theta(\mu, \Sigma) \in \Theta$, $P_{n, \theta, \frac{q}{n}}$ corresponds to the uniform distribution on $X^{(1)}, X^{(2)}, \dots, X^{(q)}$, the q points in $\mathbb{X}_n = \text{Supp}(P_n)$ at which $\|X - \mu\|_{\Sigma^{-1}}^2$ is the smallest, that is, the q -nearest neighbours of μ in \mathbb{X}_n for the Mahalanobis norm associated to Σ^{-1} .

The algorithm 3.3 adapts as follows:

Algorithm 3.5: Heteroscedastic clustering for multidimensional Gaussian mixtures

Input: \mathbb{X}_n an n -sample from P , q and k ;

Initialization

$$h = \frac{q}{n};$$

Sample $\mu_1, \mu_2, \dots, \mu_k$ from \mathbb{X}_n without replacement. ;

for i **in** $1..k$:

$$\Sigma_i = I_d;$$

$$\theta_i = \theta(\mu_i, \Sigma_i);$$

while the (μ_i, Σ_i) s vary make the following two steps:

Decomposition in weighted Voronoi cells.

for i **in** $1..k$:

$$\mathcal{C}_i = \emptyset;$$

$$c_i = P_{n, \theta_i, h} u;$$

$$x_i^2 = P_{n, \theta_i, h} \|u - c_i\|_{\Sigma_i^{-1}}^2;$$

$$d_i = \log(\det(\Sigma_i));$$

for j **in** $1..n$:

Add X_j to the \mathcal{C}_i (for i as small as possible) satisfying

```

     $\|X - c_i\|_{\Sigma_i^{-1}}^2 + x_i^2 + d_i \leq \|X - c_l\|_{\Sigma_l^{-1}}^2 + x_l^2 + d_l;$ 
# Computation of the new centres and weights.
for i in 1..k:
     $\mu_i = \frac{1}{|\mathcal{C}_i|} \sum_{X \in \mathcal{C}_i} X;$ 
     $\theta'_i = \theta(\mu_i, \Sigma_i);$ 
    for k, l in 1..d:
         $[\Sigma_i]_{k,l} = \frac{1}{|\mathcal{C}_i|} \sum_{X=(X^1, X^2, \dots, X^d) \in \mathcal{C}_i} P_{n, \theta'_i, h}(X^k - u^k)(X^l - u^l);$ 
     $\theta_i = \theta(\mu_i, \Sigma_i);$ 
Output:  $(\mu_i)_{i \in [1, k]}, (\Sigma_i)_{i \in [1, k]}, (\mathcal{C}_i)_{i \in [1, k]}$ 

```

Lemma 3.33. *In the Gaussian multivariate context, for all distributions Q and Q' on \mathbb{R}^d , a maximizer of*

$$\psi : \Sigma \mapsto \int Q Q' \langle \theta(u, \Sigma), f(v) \rangle - \psi(\theta(u, \Sigma)) du dv$$

is given by:

$$\Sigma = [\int Q Q' (v_i - u_i)(v_j - u_j) du dv]_{i, j \in [1, d]}.$$

The proof of Lemma 3.33 is to be found in Section 3.3.2.

As a consequence and as a direct corollary of Proposition 3.31, it comes.

Corollary 3.34. *Algorithm 3.5 converges to a local maximum of*

$$(\boldsymbol{\mu}, \boldsymbol{\Sigma}) \mapsto P_n \max_{i \in [1, k]} \log(p_{n, \mu_i, \Sigma_i}(u)),$$

where for every $x \in \mathbb{R}$,

$$p_{n, \mu_i, \Sigma_i}(x) = \frac{1}{\sqrt{2\pi \det(\Sigma_i)}} \exp\left(-\frac{\|x - m(P_{n, \theta_i, h})\|_{\Sigma_i^{-1}}^2 + v(P_{n, \theta_i, h})}{2}\right),$$

for $P_{n, \mu_i, h} \in \mathcal{P}_{\theta_i, h}(P_n)$, with $m(P_{n, \theta_i, h}) = P_{n, \theta_i, h} u$ and $v(P_{n, \theta_i, h}) = P_{n, \theta_i, h} \|u - m(P_{n, \theta_i, h})\|_{\Sigma_i^{-1}}^2$.

Adaptation of the algorithms to data corrupted by noise

In order to deal with outliers or clutter noise, it will be appropriate to add a trimming step in Algorithms 3.3, 3.4 and 3.5, exactly as in Algorithm 3.1. The *trimming parameter* is defined by $t = \frac{N-n}{N}$ when n is the number of points assigned to a cluster by the algorithm among the dataset of size N . Morally, a proportion t of the data will be considered as noise, that is, sent to the cell \mathcal{C}_0 .

Just before the step "Computation of the new centres and weights", we add the following step, with the notation

$$l(X) = \max_{i \in [1, k]} \langle \theta(X_j, \chi_i), P_{\theta_i, h} f(u) \rangle - \psi(\theta(X_j, \chi_i)). \quad (3.17)$$

Algorithm 3.6: Trimming step

```

# Trimming step
Sort  $[l(X_1), l(X_2), \dots, l(X_N)]$  in a non-decreasing order. ;
The sorted vector is  $[l(X^{(1)}), l(X^{(2)}), \dots, l(X^{(N)})]$  ;
 $\mathcal{C}_0 \leftarrow \emptyset$  ;
for i in 1..Nt:

```


Set j the index such that $X^{(i)} \in \mathcal{C}_j$;
 $\mathcal{C}_j \leftarrow \mathcal{C}_j \setminus \{X^{(i)}\}$;
 $\mathcal{C}_0 \leftarrow \mathcal{C}_0 \cup \{X^{(i)}\}$;

3.2.3 Experiments

The method in action

We apply Algorithm 3.5 to \mathbb{X}_{300}^1 , a 250-sample from P_1 , a mixture of 3 multivariate Gaussian distributions, with additional 50 points of clutter noise, generated uniformly on $[2, 8] \times [2, 8]$. The mass parameter chosen is $h = 0.1$ for all of the experiments. The distribution P_1 is defined by

$$P_1 = 0.5\mathcal{N}(\mu_1, \Sigma_1) + 0.3\mathcal{N}(\mu_2, \Sigma_2) + 0.2\mathcal{N}(\mu_3, \Sigma_3),$$

with $\mu_1 = (5, 4)^T$, $\mu_2 = (4, 5)^T$ and $\mu_3 = (6, 5)^T$, $\Sigma_1 = \begin{pmatrix} 0.3^2 & 0 \\ 0 & 0.6^2 \end{pmatrix}$, $\Sigma_2 = \begin{pmatrix} 0.8^2 & 0 \\ 0 & 0.4^2 \end{pmatrix}$ and $\Sigma_3 = \begin{pmatrix} 0.1^2 & 0.1 \times 0.6 \\ 0.1 \times 0.6 & 0.6^2 \end{pmatrix}$.

We proceed like for Section 3.1.5 in order to select the trimming parameter. That is, we plot the mean cost $-\frac{1}{n} \sum_{X \in \mathbb{X} \setminus \mathcal{C}_0} l(X)$ with l defined by (3.17), with respect to the trimming parameter, or equivalently, with respect to the number of elements that are assigned to a cluster in the procedure. We shall select the parameter at which there is a slight modification, like a slope increasing, on the observed curve. We selected as a parameter, the abscissa of the blue point in Figure 3.36. For each trimming parameter, or for each number of points that are to be clustered in $\llbracket 1, 300 \rrbracket$, we applied Algorithm 3.5 to the same data set \mathbb{X}_{300}^1 10 times independently, and we keep the best clustering among the 10 trials, best in the sense that the mean cost is minimised.




Figure 3.36: Mean cost as a function of the trimming parameter – noisy sample from P_1

As noted in Section 3.1.5, the quality of a clustering method may be measured in terms of the normalised mutual information (nMI), as defined in [SJ02]. In Figure 3.37, we represented the curve of the nMI as a function of the trimming parameter, where nMI refers to the normalised mutual information between the true clustering where a point has label $i \in \llbracket 1, 3 \rrbracket$ when it was sampled from $\mathcal{N}(\mu_i, \Sigma_i)$ and 0 if it was generated as clutter noise, and the clustering provided by the algorithm. Again, the algorithm was applied 10 times to the dataset \mathbb{X}_{300} and we keep the best clustering among 10 trials. The normalised mutual information has to be compared with 0.309 which corresponds to the normalised mutual information associated with the clustering in

Figure 3.35. It is also to be compared to 1 that corresponds to the normalised mutual information between two equal clusterings. Also, we can see that the maximum is attained at a trimming parameter corresponding to the slope variation in Figure 3.36. It corresponds to a configuration where the number of points removed is close to 50, the actual amount of points generated from clutter noise. The number of points removed is actually slightly inferior to the amount of points generated as noise, since some of the points are too close to some mean c_1, c_2 or c_3 .




Figure 3.37: Normalised mutual information as a function of the trimming parameter – noisy sample from P_1

The clustering associated to the trimming parameter $\frac{300-267}{300}$, marked with a blue point in Figure 3.36 and Figure 3.37 is available in Figure 3.38.




Figure 3.38: Clustering – noisy sample from P_1

We made exactly the same experiment with the distribution P_2 defined by

$$P_2 = 0.3\mathcal{N}(\mu_1, \Sigma_1) + 0.3\mathcal{N}(\mu_2, \Sigma_2) + 0.4\mathcal{N}(\mu_3, \Sigma_3),$$

with $\mu_1 = (3, 3)^T$, $\mu_2 = (0, 1)^T$ and $\mu_3 = (3, 0)^T$, $\Sigma_1 = \begin{pmatrix} 1^2 & 0 \\ 0 & 1^2 \end{pmatrix}$, $\Sigma_2 = \begin{pmatrix} 0.3^2 & 0 \\ 0 & 2^2 \end{pmatrix}$ and $\Sigma_3 = \begin{pmatrix} 2^2 & 2 \times 0.5 \\ 2 \times 0.5 & 0.5^2 \end{pmatrix}$.

In Figure 3.40, we represented the cost associated with the clustering of a 250-sample from P_2 with additional 50 points of noise, as a function of the trimming parameter. In Figure 3.40 we represented the normalised mutual information between the clustering obtained with the algorithm and the true clustering, as a function of the trimming parameter. Again, the trimming

parameter at which the mutual information is maximal corresponds to the trimming parameter at which we observe a slight deviation in the slope of the mean cost curve. The remaining clustering is plotted in Figure 3.41.




Figure 3.39: Cost as a function of the trimming parameter – noisy sample from P_2




Figure 3.40: Normalised mutual information as a function of the trimming parameter – noisy sample from P_2

Comparison to the tclust algorithm

We compared our method to the method of [GEGMMI08] with the algorithm tclust available in the R package tclust, [FGEMI12]. For the sample from P_2 , the two methods provide similar results in terms of the normalised mutual information, as enhanced by Figure 3.40 and Figure 3.42. Although the maximal normalised information with the hclust method, 0.7232, is a slightly larger than the maximal normalised information with our method, 0.7141. In Figure 3.43, we plotted the clustering associated with the method of [GEGMMI08].

Actually, the algorithm tclust fails recovering the clusters when the components of the mixtures are overlapping as enhanced by the Figure 3.47 where 250 data points were generated from the measure P_3 and 50 points generated uniformly on $[-10, 10] \times [-10, 10]$, where P_3 is defined by:

$$P_3 = 0.3\mathcal{N}(\mu_1, \Sigma_1) + 0.3\mathcal{N}(\mu_2, \Sigma_2) + 0.4\mathcal{N}(\mu_3, \Sigma_3),$$




Figure 3.41: Clustering – noisy sample from P_2

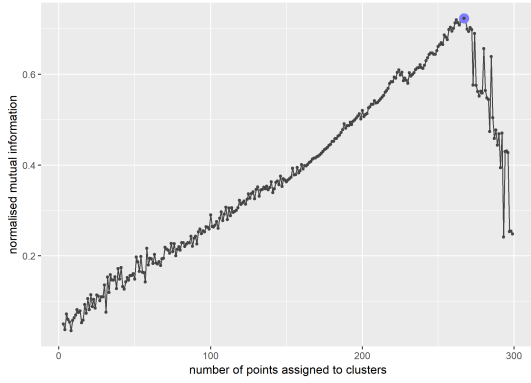


Figure 3.42: Normalised mutual information



Figure 3.43: Clustering

Figure 3.44: The tclust algorithm on the sample from P_2

with $\mu_1 = (1, 3)^T$, $\mu_2 = (0, 1)^T$ and $\mu_3 = (3, 0)^T$, $\Sigma_1 = \begin{pmatrix} 3^2 & 0 \\ 0 & 0.3^2 \end{pmatrix}$, $\Sigma_2 = \begin{pmatrix} 0.2^2 & 0 \\ 0 & 3^2 \end{pmatrix}$ and $\Sigma_3 = \begin{pmatrix} 2^2 & 2 \times 1 \\ 2 \times 1 & 1^2 \end{pmatrix}$.




Figure 3.45: tclust method




Figure 3.46: our method

Figure 3.47: Comparison to the tclust algorithm for overlapping mixtures.

3.3 Proofs

3.3.1 Proofs for Section 3.1

Proof of Lemma 3.5

For $u \in [0, 1]$, set $F_{\mathbf{c}}^{-1}(u) = r_{\phi, u}^2(\mathbf{c})$ the u -quantile of the random variable $d_{\phi}(X, \mathbf{c})$ for $X \sim P$. That is,

$$F_{\mathbf{c}}^{-1}(u) = \inf \{ r \geq 0 \mid \mathbb{P}_X(d_{\phi}(X, \mathbf{c}) \leq r^2) > u \}.$$

With the notation $\tilde{F}_{\mathbf{c}}^{*-1}(u)$ for the u -quantile of $d_{\phi}(\tilde{X}^*, \mathbf{c})$ for $\tilde{X}^* \sim P_{\mathbf{c}, h} \in \mathcal{P}_{\phi, \mathbf{c}, h}(P)$, it holds $\tilde{F}_{\mathbf{c}}^{*-1}(u) = F_{\mathbf{c}}^{-1}(hu)$. Moreover, for U a random variable uniform on $[0, 1]$, $F_{\mathbf{c}}^{*-1}(U)$ and

$d_\phi(\tilde{X}^*, \mathbf{c})$ have the same distribution. Thus, we may write:

$$\begin{aligned} V_{\phi,h}^P(\mathbf{c}) &= \mathbb{E}_{\tilde{X}^*} d_\phi(\tilde{X}^*, \mathbf{c}) \\ &= \int_{u=0}^1 F_{\mathbf{c}}^{-1}(hu) du \end{aligned}$$

Let $\tilde{P} \in \mathcal{P}_h(P)$ be a Borel probability measure on Ω such that $h\tilde{P}$ is a sub-measure of P , and let $\tilde{F}_{\mathbf{c}}^{-1}(u)$ denote the u -quantile of $d_\phi(\tilde{X}, \mathbf{c})$ for $\tilde{X} \sim \tilde{P}$. It holds that $\tilde{F}_{\mathbf{c}}^{-1}(u) \geq F_{\mathbf{c}}^{-1}(hu)$ since $P(\mathcal{B}_\phi(\mathbf{c}, r)) \geq h\tilde{P}(\mathcal{B}_\phi(\mathbf{c}, r))$. Thus, we may write:

$$\begin{aligned} V_{\phi,h}^P(\mathbf{c}) &\leq \int_{u=0}^1 \tilde{F}_{\mathbf{c}}^{-1}(u) du \\ &= \tilde{P} d_\phi(u, \mathbf{c}). \end{aligned}$$

Note that equality holds if and only if $\tilde{F}_{\mathbf{c}}^{-1}(u) = F_{\mathbf{c}}^{-1}(hu)$ for almost all $u \in [0, 1]$. Thus, $d_\phi(\tilde{X}, \mathbf{c})$ and $d_\phi(\tilde{X}^*, \mathbf{c})$ might have the same distribution. It yields that \tilde{P} is supported on $\overline{\mathcal{B}_\phi(\mathbf{c}, r_{\phi,h}(\mathbf{c}))}$ and is such that $h\tilde{P}(\mathcal{B}_\phi(\mathbf{c}, r_{\phi,h}(\mathbf{c}))) = P(\mathcal{B}_\phi(\mathbf{c}, r_{\phi,h}(\mathbf{c})))$. Since $h\tilde{P}$ is a sub-measure of P , it holds that $\tilde{P} \in \mathcal{P}_{\phi,\mathbf{c},h}(P)$. Reciprocally, equality holds for every measure in $\mathcal{P}_{\phi,\mathbf{c},h}(P)$.

Proof of Lemma 3.6

Set $0 < h < h' < 1$. With $F_{\mathbf{c}}^{-1}(u) = r_{\phi,u}^2(\mathbf{c})$ the u -quantile of the random variable $d_\phi(X, \mathbf{c})$ for $X \sim P$ and $u \in [0, 1]$, as mentioned in the proof of Lemma 3.5, we get that:

$$\begin{aligned} V_{\phi,h}^P(\mathbf{c}) &= \int_{u=0}^1 F_{\mathbf{c}}^{-1}(hu) du \\ &\leq \int_{u=0}^1 F_{\mathbf{c}}^{-1}(h'u) du \\ &= V_{\phi,h'}^P(\mathbf{c}) \end{aligned}$$

since $F_{\mathbf{c}}^{-1}$ is non-decreasing. Thus, equality holds if and only if $F_{\mathbf{c}}^{-1}(hu) = F_{\mathbf{c}}^{-1}(h'u)$ for almost all $u \in [0, 1]$. Since $F_{\mathbf{c}}^{-1}$ is non-decreasing and right-continuous, it yields that equality holds if and only if, for all $l < h'$, $F_{\mathbf{c}}^{-1}(l) = F_{\mathbf{c}}^{-1}(0)$. That is, $r_{\phi,l}^2(\mathbf{c}) = r_{\phi,0}^2(\mathbf{c})$. From (3.1) and Definition 3.3 it follows that $P(\mathcal{B}_\phi(\mathbf{c}, r_{\phi,h'}(\mathbf{c}))) = 0$. Reciprocally, equality holds when $P(\mathcal{B}_\phi(\mathbf{c}, r_{\phi,h'}(\mathbf{c}))) = 0$.

Proof of Theorem 3.9

The proof of Theorem 3.9 is based on the following lemmas.

Lemma 3.35. *Assume that ϕ is \mathcal{C}^2 . Then for every $h \in (0, 1)$ and $K > 0$, there exists $r^+ < \infty$ such that for $F_0 = \overline{\text{Conv}(\text{Supp}(P))} \subset \Omega^\circ$,*

$$\sup_{c \in F_0 \cap \overline{\mathcal{B}}(0, K)} r_{\phi,h}(c) \leq r^+.$$

As a consequence, if \mathbf{c} is a codebook with a codepoint $c_{j_0} \in F_0$ satisfying $\|c_{j_0}\| \leq K$, then

$$r_{\phi,h}(\mathbf{c}) \leq r^+.$$

Proof of Lemma 3.35. Since P is a probability measure, $P(\mathcal{B}(0, K_+)) > h$ for some $K_+ > 0$. Thus, if $c \in \overline{\mathcal{B}}(0, K)$, $P(\mathcal{B}(c, K + K_+)) > h$. Since $\mathcal{B}(c, K + K_+) \subset \mathcal{B}(0, 2K + K_+)$, and ϕ is \mathcal{C}^2 , according to the mean value theorem, there exists C_+ such that,

$$\forall x, y \in \mathcal{B}(c, K + K_+) \cap F_0, d_\phi(x, y) \leq C_+ \|x - y\|.$$

Thus, for every $c \in \overline{\mathcal{B}}(0, K)$, $P\left(\mathcal{B}_\phi\left(c, \sqrt{C_+(2K + K_+)}\right)\right) > h$.

Hence $r_{\phi,h}(c) \leq \sqrt{C_+(2K + K_+)} = r^+$.

At last, if \mathbf{c} is such that $c_{j_0} \in \overline{\mathcal{B}}(0, K) \cap F_0$, then $\mathcal{B}_\phi(c_{j_0}, r_{\phi,h}(c_{j_0})) \subset \mathcal{B}_\phi(\mathbf{c}, r_{\phi,h}(c_{j_0}))$. Therefore $P(\mathcal{B}_\phi(\mathbf{c}, r_{\phi,h}(c_{j_0}))) > h$, thus $r_{\phi,h}(\mathbf{c}) \leq r^+$. \square

Lemma 3.36. Assume that $F_0 = \overline{\text{Conv}(\text{Supp}(P))} \subset \Omega^\circ$ and ϕ is \mathcal{C}^2 on Ω . Then, for every $K > 0$, there exists $C_K > 0$ such that for every \mathbf{c} and \mathbf{c}' in $\overline{\mathcal{B}}(0, K) \cap F_0$, and $x \in \Omega$,

$$|d_\phi(x, \mathbf{c}) - d_\phi(x, \mathbf{c}')| \leq C_K \text{dist}(\mathbf{c}, \mathbf{c}') (1 + \|x\|).$$

Where we recall that $\text{dist}(\mathbf{c}, \mathbf{c}') = \min_{\sigma \in \Sigma_k} \max_{i \in [1, k]} |c_i - c'_{\sigma(i)}|$ with Σ_k the set of all permutations of $[1, k]$.

Proof of Lemma 3.36. The set $F_0 \cap \overline{\mathcal{B}}(0, K)$ is a convex compact subset of Ω° . Let $x \in \mathbb{R}^d$ and $\mathbf{c}, \mathbf{c}' \in F_0 \cap \overline{\mathcal{B}}(0, K)$. Then, for every $j \in [1, k]$, we may write

$$\begin{aligned} |d_\phi(x, c_j) - d_\phi(x, c'_j)| &\leq |\phi(c'_j) - \phi(c_j)| + \left| \langle x, \nabla_{c'_j} \phi - \nabla_{c_j} \phi \rangle \right| + \left| \langle \nabla_{c'_j} \phi, c'_j \rangle - \langle \nabla_{c_j} \phi, c_j \rangle \right| \\ &\leq C_K \|c_j - c'_j\| (1 + \|x\|), \end{aligned}$$

for some constant C_K . For this, we applied the mean value theorem to the \mathcal{C}^1 functions ϕ and $x \mapsto \nabla \phi(x)$ defined and finite on the convex compact set $F_0 \cap \overline{\mathcal{B}}(0, K)$. Thus,

$$|d_\phi(x, \mathbf{c}) - d_\phi(x, \mathbf{c}')| \leq C_K (1 + \|x\|) \max_j \|c_j - c'_j\|$$

and

$$|d_\phi(x, \mathbf{c}) - d_\phi(x, \mathbf{c}')| \leq C_K (1 + \|x\|) \text{dist}(\mathbf{c}, \mathbf{c}').$$

\square

Lemma 3.37. Assume that $F_0 = \overline{\text{Conv}(\text{Supp}(P))} \subset \Omega^\circ$, $P\|u\| < \infty$ and ϕ is \mathcal{C}^2 on Ω . The map $(s, \mathbf{c}) \rightarrow sV_{\phi,s}^P(\mathbf{c}) = R_s(\mathbf{c})$ is continuous. Moreover, for every $h \in (0, 1)$, $\epsilon > 0$ and $K > 0$, there is $s < h$ for which

$$\sup_{\mathbf{c} \in (F_0 \cap \overline{\mathcal{B}}(0, K))^{(k)}} R_h(\mathbf{c}) - R_s(\mathbf{c}) \leq \epsilon.$$

Proof of Lemma 3.37. According to Lemma 3.5 and Lemma 3.36, for every $h \in (0, 1)$,

$$\begin{aligned} V_{\phi,h}^P(\mathbf{c}) - V_{\phi,h}^P(\mathbf{c}') &\leq P_{\mathbf{c}',h} d_\phi(u, \mathbf{c}) - P_{\mathbf{c}',h} d_\phi(u, \mathbf{c}') \\ &\leq P_{\mathbf{c}',h} |d_\phi(u, \mathbf{c}) - d_\phi(u, \mathbf{c}')| \\ &\leq \frac{C_K}{h} \text{dist}(\mathbf{c}, \mathbf{c}') (1 + P\|u\|), \end{aligned}$$

for some $C_K > 0$. As a consequence, $|hV_{\phi,h}^P(\mathbf{c}) - hV_{\phi,h}^P(\mathbf{c}')| \rightarrow 0$ when $\text{dist}(\mathbf{c}, \mathbf{c}') \rightarrow 0$. Now, note that for every $s < h$,

$$\begin{aligned} R_h(\mathbf{c}) - R_s(\mathbf{c}) &= P d_\phi(u, \mathbf{c}) \left(\mathbb{1}_{\mathcal{B}_\phi(\mathbf{c}, r_{\phi,h}(\mathbf{c}))}(u) - \mathbb{1}_{\mathcal{B}_\phi(\mathbf{c}, r_{\phi,s}(\mathbf{c}))}(u) \right) \\ &\leq r_{\phi,h}(\mathbf{c}) (h - s) \end{aligned}$$

Moreover, according to Lemma 3.35, $\sup_{\mathbf{c} \in (F_0 \cap \overline{\mathcal{B}}(0, K))^{(k)}} r_{\phi,h}(\mathbf{c}) \leq r^+$ for some $r^+ < \infty$. That concludes. \square

Lemma 3.38. For every $\mathbf{c} \in \Omega^{(k)}$ there exists $j \in \llbracket 1, k \rrbracket$ such that

$$\|T(\mathbf{c})_j\| \leq \frac{kM_1}{h},$$

with $M_1 = P\|u\|$.

Proof of Lemma 3.38. Since $h \sum_{i=1}^k P_{\mathbf{c},h}^{VC,i,\mathbf{c}} = hP_{\mathbf{c},h}$ is a sub-measure of P with P -mass h , there exists $i \in \llbracket 1, k \rrbracket$ for which $hP_{\mathbf{c},h}^{VC,i,\mathbf{c}}(\Omega) \geq \frac{h}{k}$. As a consequence,

$$\|T(\mathbf{c})_i\| \leq \frac{k}{h} P\|u\|.$$

□

In the following, we will use the more concise notation $R_{k,h}^*$ for the optimal risk $hV_{\phi,k,h}(P)$ and $R_h(\mathbf{c})$ for the risk $hV_{\phi,h}^P(\mathbf{c})$ at $\mathbf{c} \in \Omega^{(k)}$.

We start with the case $k = 1$. If $\hat{\mathbf{c}}_n$ is a sequence of $\frac{1}{n}$ minimisers of $R_{1,h}^*$ in $\Omega^{(1)}$, then according to Lemma 3.38, $|T(\hat{\mathbf{c}}_n)| \leq \frac{M_1}{h}$ and by definition, $T(\hat{\mathbf{c}}_n) \in F_0$. Since $F_0 \cap \overline{\mathcal{B}}(0, \frac{M_1}{h})$ is compact, up to a subsequence, $T(\hat{\mathbf{c}}_n)$ converges to some $\mathbf{c}^* \in F_0 \cap \overline{\mathcal{B}}(0, \frac{M_1}{h})$. From Proposition 3.11 and Lemma 3.37, it comes that $R_h(\mathbf{c}^*) = R_{1,h}^*$ and Theorem 3.9 is true for $k = 1$.

Note that if $R_{k-1,h}^* - R_{k,h}^* = 0$, then there exists A with $P(A) \geq h$ such that $P|_A$ is supported on $k - 1$ points at most. Then any codebook made of these (at most) $k - 1$ points completed with arbitrary points in Ω are minimisers of $R_{k,h}^*$.

In order to complete the proof of Theorem 3.9, it remains to prove the existence of minimisers under the assumption $R_{k-1,h}^* - R_{k,h}^* > 0$.

We will prove by induction the following Lemma.

Lemma 3.39. For every $k \geq 2$, if $R_{k-1,h}^* - R_{k,h}^* > 0$, then

$$\alpha := \min_{j=2,\dots,k} R_{j-1,h}^* - R_{j,h}^* > 0.$$

Moreover there exists $0 < h^- < h < h^+ < 1$ and C_{h^-,h^+} such that, for every $j \in \llbracket 2, k \rrbracket$ and $s \in [h^-, h^+]$,

- $R_{j-1,h^-}^* - R_{j,h^+}^* \geq \frac{\alpha}{2}$.
- For every $\frac{\alpha}{4}$ -minimizer $\mathbf{c}_{j,s}^*$ of $R_{j,s}^*$, $\sup_{p \in \llbracket 1, j \rrbracket} \|T_s(\mathbf{c}_{j,s}^*)_p\| \leq C_{h^-,h^+}$.
- There is a minimizer $\mathbf{c}_{j,s}^*$ of $R_{j,s}^*$ such that $\forall p \in \llbracket 1, j \rrbracket$, $\|\mathbf{c}_{j,s,p}^*\| \leq C_{h^-,h^+}$.

Proof of Lemma 3.39. The third point follows on from the second point. Indeed, for every sequence $\mathbf{c}_{k,s}^{*(n)}$ of $\frac{\alpha}{4n}$ minimisers of $R_{k,s}^*$, for every $i \in \llbracket 1, k \rrbracket$, $\|T_s(\mathbf{c}_{k,s}^{*(n)})_i\| \leq C_{h^-,h^+}$. Since $(\overline{\mathcal{B}}(0, C_{h^-,h^+}) \cap F_0)^{(k)}$ is a compact set, the limit in $(\overline{\mathcal{B}}(0, C_{h^-,h^+}) \cap F_0)^{(k)}$ of any converging subsequence of $(T_s(\mathbf{c}_{k,s}^{*(n)}))_n$ is a minimiser of $R_{k,s}^*$.

First note that if there exists $j \leq k$ such that $R_{j-1,h}^* - R_{j,h}^* = 0$, then there will exist a set A with $P(A) \geq h$ such that $P|_A$ will be supported on $j - 1$ points. Thus it will come that $R_{k-1,h}^* = R_{k,h}^* = 0$. As a consequence, when $R_{k-1,h}^* - R_{k,h}^* > 0$, α is positive. In the following we make this assumption.

In order to prove the other points, we proceed recursively. Assume that $k = 2$, then for every $s > 0$, according to Lemma 3.38 and Proposition 3.11, we can take $\mathbf{c}_{1,s}^*$ a $R_{1,s}^*$ minimiser that satisfies

$$\|\mathbf{c}_{1,s}^*\| \leq \frac{P\|u\|}{s}.$$

Denote by \mathbf{c}_{1,h^-}^* a minimizer of R_{1,h^-}^* , and $\mathbf{c}_{2,h}^*$ an $\frac{\alpha}{8}$ minimizer of $R_{2,h}^*$. According to Lemma 3.37, for a fixed \mathbf{c} , $s \mapsto R_s(\mathbf{c})$ is continuous, thus we may choose h^+ such that $R_{h^+}(\mathbf{c}_{2,h}^*) \leq R_h(\mathbf{c}_{2,h}^*) + \frac{\alpha}{8}$. Then,

$$\begin{aligned} R_{2,h^+}^* &\leq R_{h^+}(\mathbf{c}_{2,h}^*) \\ &\leq R_h(\mathbf{c}_{2,h}^*) + \frac{\alpha}{8} \\ &\leq R_{2,h}^* + \frac{\alpha}{4}. \end{aligned}$$

On the other hand, set $h^{--} = \frac{h}{2}$. Then $\sup_{s \geq h^{--}} \|\mathbf{c}_{1,s}^*\| \leq \frac{P\|u\|}{h^{--}} = C_{h^{--}}$. According to Lemma 3.37, there exists $h > h^- \geq h^{--}$ such that $\sup_{\|\mathbf{c}\| \leq C_{h^{--}}} (R_h(\mathbf{c}) - R_{h^-}(\mathbf{c})) \leq \frac{\alpha}{4}$. For such an h^- , we may write

$$\begin{aligned} R_{1,h^-}^* &= R_{h^-}(\mathbf{c}_{1,h^-}^*) \\ &\geq R_h(\mathbf{c}_{1,h^-}^*) - \frac{\alpha}{4} \\ &\geq R_{1,h}^* - \frac{\alpha}{4}. \end{aligned}$$

Since $R_{k,h}^* - R_{k-1,h}^* \geq \alpha$, it comes that $R_{1,h^-}^* - R_{2,h^+}^* \geq \frac{\alpha}{2}$.

Now, if \mathbf{c} is an $\alpha/4$ -minimizer of $R_{2,s}^*$, for $s \geq h - (h - h^-)/2 := h_b^-$, its Bregman-Voronoi cells restricted to the h -trimming set, $V_{j,s}$, have weight not smaller than $h - h_b^-$. Indeed, suppose that $P(V_{1,s}) < h - h_b^-$. Then

$$\begin{aligned} R_{2,h^+}^* &\geq R_s(\mathbf{c}) - \frac{\alpha}{4} \\ &\geq \text{Pd}_\phi(u, c_2) \mathbb{1}_{V_{2,s}}(u) - \frac{\alpha}{4} \\ &\geq R_{1,h_b^-}^* - \frac{\alpha}{4} \\ &\geq R_{1,h^-}^* - \frac{\alpha}{4}, \end{aligned}$$

hence the contradiction. Choosing $h^- = h_b^-$ entails that $\|T_s(c_{j,s}^*)\| \leq \frac{P\|u\|}{h-h^-}$, this gives the result for $k = 2$.

Assume that the proposition is true for index $k - 1$, we will prove that it is also true for index k . Set $\alpha = \min_{j \in [2,k]} R_{j-1,h}^* - R_{j,h}^* > 0$. Set h^{--} and h^{++} the elements h^- and h^+ associated with step $k - 1$. Set \mathbf{c}_{k-1,h^-}^* a minimiser of R_{k-1,h^-}^* and $\mathbf{c}_{k,h}^*$ an $\frac{\alpha}{8}$ -minimiser of $R_{k,h}^*$. For $h < h^+ < h^{++}$ such that $R_{h^+}(\mathbf{c}_{k,h}^*) \leq R_h(\mathbf{c}_{k,h}^*) + \frac{\alpha}{8}$ that exists according to Lemma 3.37. It comes that $R_{k,h^+}^* \leq R_{k,h}^* + \frac{\alpha}{4}$.

On the other hand, according to Lemma 3.37, it comes that

$$\sup_{\mathbf{c} \in (\overline{\mathcal{B}}(0, C_{h^{--}, h^{++}}) \cap F_0)^{(k)}} (R_h(\mathbf{c}) - R_{h^-}(\mathbf{c})) \leq \frac{\alpha}{4}$$

for some $h > h^- > h^{--}$.

Then, according to step $k-1$ and Proposition 3.11, $T_{h^-}(\mathbf{c}_{k-1,h^-}^*)$ is a minimiser of R_{k-1,h^-}^* in $(\overline{\mathcal{B}}(0, C_{h^{--}, h^{++}}) \cap F_0)^{k-1}$, thus we may write

$$R_{k-1,h^-}^* = R_{h^-}(T_{h^-}(\mathbf{c}_{k-1,h^-}^*)) \geq R_{k-1,h}^* - \frac{\alpha}{4}.$$

As a consequence, and since $R_{k-1,h}^* - R_{k,h}^* \geq \alpha$, it comes that $R_{k-1,h^-}^* - R_{k,h^+}^* \geq \frac{\alpha}{2}$. Now, if \mathbf{c} is an $\frac{\alpha}{4}$ -minimiser of $R_{k,s}^*$, for $s \geq h_b^- = \frac{h+h^-}{2}$, if for instance $P(V_{1,s}) < h - h_b^-$, it comes that

$$\begin{aligned} R_{k,h^+}^* &\geq R_s(\mathbf{c}) - \frac{\alpha}{4} \\ &\geq P \sum_{j=2}^k d_\phi(u, c_j) \mathbb{1}_{V_{j,s}}(u) - \frac{\alpha}{4} \\ &\geq R_{k-1,h_b^-}^* - \frac{\alpha}{4} \\ &\geq R_{k-1,h^-}^* - \frac{\alpha}{4}, \end{aligned}$$

which is a contradiction. Thus, with the choice $h^- = h_b^-$, $P(V_{p,s}) \geq h - h^-$ for every $p \in \llbracket 1, k \rrbracket$, and the lemma is true for index k . \square

Proof of Proposition 3.11

For $P_{\mathbf{c},h} \in \mathcal{P}_{\phi,\mathbf{c},h}(P)$ and $i \in \llbracket 1, k \rrbracket$, set $T(\mathbf{c})_i = m(P_{\mathbf{c},h}^{BV,i,\mathbf{c}})$. Then,

$$\begin{aligned} V_{\phi,h}^P(\mathbf{c}) &= P_{\mathbf{c},h} d_\phi(u, \mathbf{c}) \\ &= \sum_{i=1}^k P_{\mathbf{c},h}^{BV,i,\mathbf{c}}(\Omega) \frac{P_{\mathbf{c},h}^{BV,i,\mathbf{c}}}{P_{\mathbf{c},h}^{BV,i,\mathbf{c}}(\Omega)} d_\phi(u, c_i) \\ &\geq \sum_{i=1}^k P_{\mathbf{c},h}^{BV,i,\mathbf{c}}(\Omega) \frac{P_{\mathbf{c},h}^{BV,i,\mathbf{c}}}{P_{\mathbf{c},h}^{BV,i,\mathbf{c}}(\Omega)} d_\phi(u, T(\mathbf{c})_i) \end{aligned}$$

according to Lemma 3.2. Moreover, provided that $P_{\mathbf{c},h}^{BV,i,\mathbf{c}}(\Omega) \neq 0$, equality holds if and only if $c_i = T(\mathbf{c})_i$. Finally, according to Lemma 3.5,

$$V_{\phi,h}^P(\mathbf{c}) \geq P_{\mathbf{c},h} d_\phi(u, T(\mathbf{c})) \geq P_{T(\mathbf{c}),h} d_\phi(u, T(\mathbf{c})) = V_{\phi,h}^P(T(\mathbf{c})).$$

If \mathbf{c}^* is a minimizer of $\mathbf{c} \mapsto V_{\phi,h}^P(\mathbf{c}) = P_{\mathbf{c},h} d_\phi(u, \mathbf{c})$ for any $P_{\mathbf{c},h} \in \mathcal{P}_{\phi,\mathbf{c},h}(P)$. From the definition of \mathbf{c}^* , it comes that inequalities are equalities. Thus, $c_i^* = T(\mathbf{c}^*)_i$ which concludes.

Proof of Lemma 3.12

Set $P_{\mathbf{c},h} \in \mathcal{P}_{\phi,\mathbf{c},h}(P)$ and $P_{\mathbf{c}',h} \in \mathcal{P}_{\phi,\mathbf{c}',h}(P)$, then,

$$\begin{aligned} V_{\phi,h}^P(\mathbf{c}') &= P_{\mathbf{c}',h} d_\phi(u, \mathbf{c}') \\ &\leq P_{\mathbf{c},h} d_\phi(u, \mathbf{c}') \end{aligned}$$

thanks to Lemma 3.5. Thus,

$$\begin{aligned} V_{\phi,h}^P(\mathbf{c}') &\leq P_{\mathbf{c},h} \min_{i \in \llbracket 1, k \rrbracket} \phi(u) - \phi(c'_i) - \langle \nabla_{c'_i} \phi, u - c'_i \rangle \\ &\leq P_{\mathbf{c},h} \min_{i \in \llbracket 1, k \rrbracket} (\phi(u) - \phi(c_i) - \langle \nabla_{c_i} \phi, u - c_i \rangle) + \left(\phi(c_i) - \phi(c'_i) - \langle \nabla_{c'_i} \phi, u - c'_i \rangle + \langle \nabla_{c_i} \phi, u - c_i \rangle \right) \end{aligned}$$

Thus,

$$\begin{aligned} V_{\phi,h}^P(\mathbf{c}') &= P_{\mathbf{c},h} \min_{i \in \llbracket 1, k \rrbracket} d_\phi(u, c_i) + \langle \nabla_{c_i} \phi - \nabla_{c'_i} \phi, u - c'_i \rangle - d_\phi(c'_i, c_i) \\ &\leq \sum_{i=1}^k P_{\mathbf{c},h}^{BV,i,\mathbf{c}} d_\phi(u, c_i) + P_{\mathbf{c},h}^{BV,i,\mathbf{c}}(\Omega) \left(\langle \nabla_{c_i} \phi - \nabla_{c'_i} \phi, m \left(P_{\mathbf{c},h}^{BV,i,\mathbf{c}} \right) - c'_i \rangle - d_\phi(c'_i, c_i) \right) \\ &= V_{\phi,h}^P(\mathbf{c}) + \sum_{i=1}^k P_{\mathbf{c},h}^{BV,i,\mathbf{c}}(\Omega) \left(\langle \nabla_{c_i} \phi - \nabla_{c'_i} \phi, m \left(P_{\mathbf{c},h}^{BV,i,\mathbf{c}} \right) - c'_i \rangle - d_\phi(c'_i, c_i) \right). \end{aligned}$$

Note that $d_\phi(c'_i, c_i) \geq 0$. Similarly, we get that

$$V_{\phi,h}^P(\mathbf{c}') \leq V_{\phi,h}^P(\mathbf{c}) + \sum_{i=1}^k P_{\mathbf{c},h}^{BV,i,\mathbf{c}}(\Omega) \left(\langle \nabla_{c_i} \phi - \nabla_{c'_i} \phi, m \left(P_{\mathbf{c},h}^{BV,i,\mathbf{c}} \right) - c_i \rangle + d_\phi(c_i, c'_i) \right).$$

Proof of Theorem 3.14

Before proving Theorem 3.14, first note that since ϕ is strictly convex and continuous, $\psi : x \mapsto \phi(x) - \langle x, a \rangle + b$ is also strictly convex and continuous, for every a and b , thus $\psi^{-1}(0)$ is a closed set. Thus, $\mathbb{1}_{\psi^{-1}(0)}$ is measurable, and after an integration by substitution, since the strict convexity of ψ entails that there are at most two points of $\psi^{-1}(0)$ in any line containing the point 0, the Lebesgue measure of $\psi^{-1}(0)$ is 0. Since P is absolutely continuous with respect to the Lebesgue measure, it follows that boundaries of Bregman balls have P -mass equal to 0.

The proof of Theorem 3.14 is based on the following lemma:

Lemma 3.40. *Set $(P_n)_{n \in \mathbb{N}}$, a sequence of probabilities that converges weakly to a distribution P . Assume that $\text{Supp}(P_n) \subset \text{Supp}(P) \subset \mathbb{R}^d$, $F_0 = \overline{\text{Conv}(\text{Supp}(P))} \subset \Omega^\circ$ and ϕ is \mathcal{C}_2 on Ω . Then, for every $h \in (0, 1)$ and $K > 0$, there exists $K_+ > 0$ such that for every $\mathbf{c} \in \Omega^{(k)}$ such that $|c_i| \leq K$ for some $i \in \llbracket 1, k \rrbracket$ and every $n \in \mathbb{N}$,*

$$r_{n,\phi,h}(\mathbf{c}) \leq r_+ = \sqrt{4(2K + K_+) \sup_{c \in F_0 \cap \overline{\mathcal{B}}(0, 2K + K_+)} \|\nabla_c \phi\|}.$$

Proof of Lemma 3.40. Set $c \in \overline{\mathcal{B}}(0, K) \cap F_0$. Since P_n converges weakly to P , according to the Prokhorov theorem, $(P_n)_n$ is tight. Thus, there is $K_+ > 0$ such that $P_n(\overline{\mathcal{B}}(0, K_+)) > h$ for all $n \in \mathbb{N}$ and $P(\overline{\mathcal{B}}(0, K_+)) > h$. It comes that $P_n(\overline{\mathcal{B}}(c, K + K_+)) > h$. Moreover, for every x, y in $F_0 \cap \overline{\mathcal{B}}(0, 2K + K_+)$, the mean value theorem yields

$$\begin{aligned} d_\phi(x, y) &\leq 2 \sup_{c \in F_0 \cap \overline{\mathcal{B}}(0, 2K + K_+)} \|\nabla_c \phi\| \|x - y\| \\ &\leq 4(2K + K_+)C_+ = (r^+)^2, \end{aligned}$$

for $C_+ = \sup_{c \in F_0 \cap \overline{\mathcal{B}}(0, 2K + K_+)} \|\nabla_c \phi\|$ that is finite since $F_0 \cap \overline{\mathcal{B}}(0, 2K + K_+)$ is compact. Thus, it comes that

$$\overline{\mathcal{B}}(c, K + K_+) \subset \overline{\mathcal{B}}_\phi(c, r_+). \tag{3.18}$$

As a consequence, $P_n(\overline{\mathcal{B}}_\phi(c, r_+)) > h$, $P_n(\overline{\mathcal{B}}_\phi(\mathbf{c}, r_+)) > h$ if $c \in \mathbf{c}$ and $r_{n,\phi,h}(\mathbf{c}) \leq r_+$. \square

The proof of Theorem 3.14 is an adaptation of the proof of Theorem 3.4 in [CAGM97].

For $(X'_n)_{n \in \mathbb{N}}$ a sequence of independent random variables from P , we can define $P_n = \frac{1}{n} \sum_{i=1}^n \delta_{X'_i}$ the empirical distribution associated with the n first realisations. Note that P_n is random.

According to Proposition 3.42, provided that $P\|u\|^p < +\infty$, for some $p > 2$, there exists $C_P > 0$ such that for some $N \in \mathbb{N}$, $\sum_{n \geq N} P(\max_{i \in \llbracket 1, k \rrbracket} |\hat{c}_{n,i}| > C_P) < \infty$. Thus, according to the Borel-Cantelli Lemma, a.s. for n large enough, for every $i \in \llbracket 1, k \rrbracket$, $|\hat{c}_{n,i}| \leq C_P$.

According to Varadarajan, in [Var58a], a.s. P_n converges weakly to P . Now we fix $(P_n)_{n \in \mathbb{N}}$, $C_P > 0$ and $N \in \mathbb{N}$ such that P_n converges weakly to P , and $\forall n \geq N, \forall i \in \llbracket 1, k \rrbracket, |\hat{c}_{n,i}| \leq C_P$. As aforementioned, this occurs with probability 1.

According to the Skorokhod's representation theorem in the Polish space \mathbb{R}^d , it is possible to construct a measured space $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{P})$ and a sequence of random variables $(X_n)_{n \in \mathbb{N}}$ and a random variable X on $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{P})$ such that $X_n \sim P_n, X \sim P$ and X_n converges to X \tilde{P} -a.s.

Set \mathbf{c}^* , the unique minimiser of $\mathbf{c} \mapsto V_{\phi,h}^P(\mathbf{c}), r'_n = r_{n,\phi,h}(\mathbf{c}^*)$ and τ'_n a $[0, 1]$ -valued measurable function such that $hP_{n,\mathbf{c}^*,h} = P\tau'_n$, that is, such that $P\tau'_n(u) = h$ and

$$\mathbb{1}_{\mathcal{B}_\phi(\mathbf{c}^*, r'_n)} \leq \tau'_n \leq \mathbb{1}_{\overline{\mathcal{B}}_\phi(\mathbf{c}^*, r'_n)}.$$

According to Lemma 3.40 with $K = |c_1^*|$ for instance, it comes for some finite r^+ that $r'_n \leq r^+$. Thus, up to a subsequence, we may assume that $r'_n \rightarrow r'_0$ for some $r'_0 \leq r_+$.

Moreover,

$$|d_\phi(X_n, \mathbf{c}^*) - d_\phi(X, \mathbf{c}^*)| \leq |\phi(X_n) - \phi(X)| + \max_{j \in \llbracket 1, k \rrbracket} \|\nabla_{c_j^*} \phi\| |X_n - X| \rightarrow 0, \quad (3.19)$$

when $n \rightarrow \infty$. As a consequence, $\tau'_n(X_n) \rightarrow \mathbb{1}_{\mathcal{B}_\phi(\mathbf{c}^*, r'_0)}(X)$ \tilde{P} -a.e. The dominated convergence theorem yields $h = P_n \tau'_n(u) \rightarrow P(\mathcal{B}_\phi(\mathbf{c}^*, r'_0))$. Thus, $\mathbb{1}_{\mathcal{B}_\phi(\mathbf{c}^*, r'_0)} = \tau_0$ P -a.e where τ_0 denotes the trimming set associated with \mathbf{c}^* and P . Moreover, since $\tau'_n(X_n) d_\phi(X_n, \mathbf{c}^*)$ is bounded by r_+ and converges to $\tau_0(X) d_\phi(X, \mathbf{c}^*)$ a.e. the dominated convergence theorem entails

$$V_{k,\phi,h}(P_n) \leq V_{\phi,h}^{P_n}(\mathbf{c}^*) \leq \frac{1}{h} \mathbb{E} [\tau'_n(X_n) d_\phi(X_n, \mathbf{c}^*)] \rightarrow \frac{1}{h} \mathbb{E} [\tau_0(X) d_\phi(X, \mathbf{c}^*)].$$

Thus, up to a subsequence,

$$\limsup_{n \rightarrow \infty} V_{k,\phi,h}(P_n) \leq V_{k,\phi,h}(P).$$

Since for $n \geq N$, for every $i \in \llbracket 1, k \rrbracket$, $|\hat{c}_{n,i}| \leq C_P$, up to a subsequence, $\hat{c}_{n,i} \rightarrow c_i$ for some $c_i \in F_0 \cap \overline{\mathcal{B}}(0, C_P)$. Set $\mathbf{c} = (c_1, c_2, \dots, c_k)$. Again, according to Lemma 3.40 with $K = C_P$, it comes that up to a subsequence, $r_{n,\phi,h}(\hat{\mathbf{c}}_n) \rightarrow r$ for some $r \geq 0$. As a consequence, from (3.19), Lemma 3.36 and the continuity of P ,

$$\lim_{n \rightarrow \infty} \tau_n(X_n) = \mathbb{1}_{\mathcal{B}_\phi(\mathbf{c}, r)}(X) a.e.$$

According to the dominated convergence theorem, it comes that $h = P(\mathcal{B}_\phi(\mathbf{c}, r)) = P_n(\tau_n(u))$.

Again, the dominated convergence theorem entails that up to a subsequence,

$$\liminf_{n \rightarrow \infty} V_{k,\phi,h}(P_n) \geq \frac{1}{h} \int \mathbb{1}_{\mathcal{B}_\phi(\mathbf{c}, r)}(u) d_\phi(u, \mathbf{c}) \geq V_{k,\phi,h}(P).$$

As a consequence, $\lim_{n \rightarrow \infty} V_{k,\phi,h}(P_n) = V_{k,\phi,h}(P)$ and $\mathbf{c} = \mathbf{c}^*$. Since every subsequence of $(\hat{\mathbf{c}}_n)_{n \in \mathbb{N}}$ has a converging subsequence to \mathbf{c}^* , the sequence $(\hat{\mathbf{c}}_n)_{n \in \mathbb{N}}$ converges to \mathbf{c}^* .

Proof of Theorem 3.15

The proof of Theorem 3.15 is based on the following results.

Lemma 3.41. *If $P\|u\|^p < \infty$ for some $p \geq 2$, then, there exists some positive constant C such that with probability larger than $1 - n^{-\frac{p}{2}}$,*

$$P_n\|u\| \leq C.$$

Proof. According to the Markov inequality,

$$\mathbb{P}(P_n\|u\| - P\|u\| \geq \epsilon) \leq \frac{\mathbb{E} \left[\left| \frac{1}{n} \sum_{i=1}^n \|X_i\| - P\|u\| \right|^p \right]}{\epsilon^p},$$

then it comes that

$$\mathbb{P}(P_n\|u\| - P\|u\| \geq \epsilon) \leq \frac{\mathbb{E} \left[\left| \sum_{i=1}^n \|X_i\| - P\|u\| \right|^p \right]}{n^p \epsilon^p},$$

and according to the Marcinkiewicz-Zygmund inequality applied to the real-valued centred random variables $Y_i = \|X_i\| - P\|u\|$ and the Minkowski inequality,

$$\begin{aligned} \mathbb{E} \left[\left| \sum_{i=1}^n \|X_i\| - P\|u\| \right|^p \right] &= \mathbb{E} \left[\left| \sum_{i=1}^n Y_i \right|^p \right] \\ &\leq B_p \mathbb{E} \left[\left(\sum_{i=1}^n Y_i^2 \right)^{\frac{p}{2}} \right] \\ &\leq B_p \left(\sum_{i=1}^n (\mathbb{E}|Y_i|^p)^{\frac{2}{p}} \right)^{\frac{p}{2}} \\ &= B_p n^{\frac{p}{2}} \mathbb{E} [|Y|^p] \\ &= B_p n^{\frac{p}{2}} P(\|u\| - P\|u\|)^p, \end{aligned}$$

for some positive constant B_p . Since $P(\|u\| - P\|u\|)^p \leq P\|u\|^p + (P\|u\|)^p \leq 2P\|u\|^p$ according to Jensen inequality, the result comes from a suitable choice of ϵ . \square

Proposition 3.42. *Assume that $P\|u\|^p < +\infty$ for some $p \geq 2$. Let h^- and h^+ denote the quantities such that $\min_{j=2,\dots,k} R_{j-1,h^-}^* - R_{j,h^+}^* \geq \frac{\alpha}{2}$ with $h^- < h < h^+$, as in Lemma 3.39. Denote by $\eta = \frac{h-h^-}{k-1}$. Then there exists C_P such that, for n large enough, with probability larger than $1 - n^{-\frac{p}{2}}$, we have, for all $j = 2, \dots, k$, and $i = 1, \dots, j$,*

$$\sup_{h-(k-j)\eta \leq s \leq h} \|\hat{c}_{j,s,i}\| \leq C_P,$$

where $\hat{c}_{j,s}$ denotes a j -codepoints empirical risk minimizer with trimming level s .

Proof. Proof of Proposition 3.42 We let C_P denote a constant to be fixed later. Similarly to the proof of Proposition 3.2, we denote by \hat{T} the operator that maps \hat{c} to the empirical means of its Bregman-Voronoi cells. Let h^- and h^+ be as in Lemma 3.39. Then, according to Lemma 3.59, for n large enough we have that,

$$\sup_{\mathbf{c} \in (\mathbb{R}^d)^{(k)}, r \geq 0} |(P - P_n)\mathcal{B}_\phi(\mathbf{c}, r)| \leq \min \left(\frac{\eta}{2}, h^+ - h \right). \tag{3.20}$$

with probability larger than $1 - \frac{1}{8n^{\frac{p}{2}}}$. On this probability event, from Lemma 3.35 and the fact that $s \mapsto r_{\phi,s}(\mathbf{c})$ is non-decreasing, we deduce that for some $r^+ > 0$,

$$\sup_{c \in \mathcal{B}(0, C_P) \cap F_0, s \leq h^+} r_{n,\psi,s}(c) \vee r_{\phi,s}(c) \leq r^+.$$

Since $P\|u\|^p < +\infty$, Lemma 3.41 yields that $P_n\|u\| \leq C_1$, for C_1 large enough, with probability larger than $1 - \frac{1}{8n^{\frac{p}{2}}}$. Besides, choosing $x = \log(8n^{\frac{p}{2}})$ in Lemma 3.60, we also have, with probability larger than $1 - \frac{1}{8n^{\frac{p}{2}}}$,

$$\sup_{c \in (\mathcal{B}(0, C_P) \cap F_0)^k, r \leq r^+} \left| (P - P_n)d_\phi(u, \mathbf{c}) \mathbb{1}_{\mathcal{B}_\phi(\mathbf{c}, r)}(u) \right| \leq \alpha_n, \quad (3.21)$$

where $\alpha_n = O(\sqrt{\log(n)/n})$. We then work on the global probability event on which all these deviation inequalities are satisfied, that have probability larger than $1 - \frac{1}{n^{\frac{p}{2}}}$, and proceed recursively on j .

For $j = 1$ and $s \geq h^-$, according to Proposition 3.11, $\hat{T}_s(\hat{\mathbf{c}}_{1,s}) = \hat{\mathbf{c}}_{1,s}$, we may write

$$\|\hat{\mathbf{c}}_{1,s}\| \leq \frac{P_n\|u\|}{h^-} \leq C_P.$$

Now assume that the statement of Proposition 3.42 holds up to order $j - 1$, let $\hat{\mathbf{c}}_{j,s}$ be a j -points empirically optimal codebook with trimming level $s \geq h - (k - j)\eta$, and assume that there exists one cell (say V_1) such that $P_n(V_1(\hat{\mathbf{c}}_{j,s}) \cap \mathcal{B}_\phi(\hat{\mathbf{c}}_{j,s}, r_{n,\phi,s}(\hat{\mathbf{c}}_{j,s})) \leq \frac{\eta}{2}$. On the one hand, we may write

$$\begin{aligned} \hat{R}_s(\hat{\mathbf{c}}_{j,s}) &\leq \hat{R}_s(\mathbf{c}_{j,h^+}^*) \\ &\leq P_n d_\phi(u, \mathbf{c}_{j,h^+}^*) \mathbb{1}_{\mathcal{B}_\phi(\mathbf{c}_{j,h^+}^*, r_{\phi,h^+}(\mathbf{c}_{j,h^+}^*))}(u) \\ &\leq R_{j,h^+}^* + \alpha_n, \end{aligned}$$

with \mathbf{c}_{j,h^+}^* a R_{j,h^+}^* minimizer, that exists according to Lemma 3.9. We can choose $C_P = \frac{C_1}{h^-} \vee C_{h^-,h^+} \vee C'_P$, where C'_P corresponds to the constant C_P for the step $j - 1$ and C_{h^-,h^+} is given by Lemma 3.39.

On the other hand, set $h' = h - (k - j + \frac{1}{2})\eta \geq h^-$, we have

$$\begin{aligned} \hat{R}_s(\hat{\mathbf{c}}_{j,s}) &\geq \sum_{p=2}^j P_n d_\phi(u, \hat{\mathbf{c}}_{j,s,p}) \mathbb{1}_{V_p(\hat{\mathbf{c}}_{j,s}) \cap \mathcal{B}_\phi(\hat{\mathbf{c}}_{j,s}, r_{n,\phi,s}(\hat{\mathbf{c}}_{j,s}))}(u) \\ &\geq \hat{R}_{h'}(\hat{\mathbf{c}}_{j-1,h'}) \\ &\geq P d_\phi(u, \hat{\mathbf{c}}_{j-1,h'}) \mathbb{1}_{\mathcal{B}_\phi(\hat{\mathbf{c}}_{j-1,h'}, r_{\phi,h'}(u))} - \alpha_n, \end{aligned}$$

according to the recursion assumption and (3.21). Thus,

$$\hat{R}_s(\hat{\mathbf{c}}_{j,s}) \geq R_{j-1,h'}^* - \alpha_n,$$

which is impossible for $\alpha_n < \frac{\alpha}{4}$. Thus, for every $p \in \llbracket 1, j \rrbracket$,

$$P_n(V_p(\hat{\mathbf{c}}_{j,s}) \cap \mathcal{B}_\phi(\hat{\mathbf{c}}_{j,s}, r_{n,\phi,s}(\hat{\mathbf{c}}_{j,s}))) \geq \frac{\eta}{2}.$$

According to Proposition 3.11, equality $\hat{T}(\hat{\mathbf{c}}_{j,s}) = \hat{\mathbf{c}}_{j,s}$ holds and entails

$$\|\hat{\mathbf{c}}_{j,s,p}\| \leq \frac{P_n\|u\|}{h - (k - j)\eta} \leq \frac{P_n\|u\|}{h^-} \leq C_P.$$

□

We assume that all the probability events described in the proof of Proposition 3.42 hold. This occurs with probability larger than $1 - n^{-\frac{p}{2}}$. On this probability event, recall that for all j $\|\hat{\mathbf{c}}_{n,j}\| \leq C_P$, $\sup_{\mathbf{c} \in (F_0 \cap \overline{\mathcal{B}}(0, C_P))^{(k)}} r_{n,\phi,h}(\mathbf{c}) \vee r_{\phi,h}(\mathbf{c}) \leq r^+$. Next we further assume that the deviation bounds of Lemma 3.60 and Lemma 3.59 hold, with parameter C_P and r_+ , to define a global probability event with mass large than $1 - n^{-\frac{p}{2}} - 2e^{-x}$. On this event, we have

$$\begin{aligned} R_h(\hat{\mathbf{c}}_n) - R_{k,h}^* &= Pd_\phi(u, \hat{\mathbf{c}}_n) \mathbb{1}_{\mathcal{B}_\phi(\hat{\mathbf{c}}_n, r_{\phi,h}(\hat{\mathbf{c}}_n))}(u) - R_{k,h}^* \\ &\leq Pd_\phi(u, \hat{\mathbf{c}}_n) \mathbb{1}_{\mathcal{B}_\phi(\hat{\mathbf{c}}_n, r_{n,\phi,h}(\hat{\mathbf{c}}_n))}(u) + Pd_\phi(u, \hat{\mathbf{c}}_n) (\mathbb{1}_{\mathcal{B}_\phi(\hat{\mathbf{c}}_n, r_{\phi,h}(\hat{\mathbf{c}}_n))}(u) - \mathbb{1}_{\mathcal{B}_\phi(\hat{\mathbf{c}}_n, r_{n,\phi,h}(\hat{\mathbf{c}}_n))}(u)) \\ &\quad - Pd_\phi(u, \mathbf{c}^*) \mathbb{1}_{\mathcal{B}_\phi(\mathbf{c}^*, r_{n,\phi,h}(\mathbf{c}^*))}(u) + Pd_\phi(u, \mathbf{c}^*) (\mathbb{1}_{\mathcal{B}_\phi(\mathbf{c}^*, r_{n,\phi,h}(\mathbf{c}^*))}(u) - \mathbb{1}_{\mathcal{B}_\phi(\mathbf{c}^*, r_{\phi,h}(\mathbf{c}^*))}(u)) \\ &\leq 2 \sup_{\mathbf{c} \in (F_0 \cap \overline{\mathcal{B}}(0, C_P))^{(k)}, r \leq r^+} |(P - P_n) d_\phi(u, \mathbf{c}) \mathbb{1}_{\mathcal{B}_\phi(\mathbf{c}, r)}(u)| + 2r^+ \sup_{\mathbf{c} \in \Omega^{(k)}, r \geq 0} |(P - P_n) \mathcal{B}_\phi(\mathbf{c}, r)| \\ &\leq \frac{C_P}{\sqrt{n}} (1 + \sqrt{x}), \end{aligned}$$

for some constant C_P .

Proof of Corollary 3.16

Denote by A the probability event described in Proposition 3.42, and by $\hat{\Delta}$ the quantity $Pd_\phi(u, \hat{\mathbf{c}}_n) \mathbb{1}_{\mathcal{B}_\phi(\hat{\mathbf{c}}_n, r_{\phi,h}(\hat{\mathbf{c}}_n))}(u) - R_{k,h}^*$. Then

$$\mathbb{E} \hat{\Delta} = \mathbb{E} \hat{\Delta} \mathbb{1}_A + \mathbb{E} \hat{\Delta} \mathbb{1}_{A^c}.$$

As in the proof of Theorem 3.15,

$$\begin{aligned} \mathbb{E} \hat{\Delta} \mathbb{1}_A &\leq \mathbb{E} \left[2 \sup_{\mathbf{c} \in (F_0 \cap \overline{\mathcal{B}}(0, C_P))^{(k)}, r \leq r^+} |(P - P_n) d_\phi(u, \mathbf{c}) \mathbb{1}_{\mathcal{B}_\phi(\mathbf{c}, r)}(u)| + 2r^+ \sup_{\mathbf{c} \in \Omega^{(k)}, r \geq 0} |(P - P_n) \mathcal{B}_\phi(\mathbf{c}, r)| \right] \\ &\leq \frac{C_{P,k,d}}{\sqrt{n}}, \end{aligned}$$

according to the proof of Lemma 3.60. It remains to bound $\mathbb{E} \hat{\Delta} \mathbb{1}_{A^c}$. For short assume that $c_0 = 0$, and the $nP_n(V_1(\hat{\mathbf{c}}_n) \cap \mathcal{B}_\phi(\hat{\mathbf{c}}_n, r_{n,\phi,h}(\hat{\mathbf{c}}_n))) \geq \frac{h}{k}$. It is straightforward that

$$Pd_\phi(u, \hat{\mathbf{c}}_n) \mathbb{1}_{\mathcal{B}_\phi(\hat{\mathbf{c}}_n, r_{\phi,h}(\hat{\mathbf{c}}_n))}(u) \leq Pd_\phi(u, \hat{\mathbf{c}}_1) \mathbb{1}_{\mathcal{B}_\phi(\hat{\mathbf{c}}_1, r_{\phi,h}(\hat{\mathbf{c}}_1))}(u).$$

Moreover, $\|\hat{\mathbf{c}}_1\| \leq \frac{k}{nh} \sum_{i=1}^n \|X_i\|$. Thus,

$$\begin{aligned} \mathbb{E} \hat{\Delta} \mathbb{1}_{A^c} &\leq \mathbb{E} [P\phi(u) - \phi(\hat{\mathbf{c}}_1) - \langle \nabla_{\hat{\mathbf{c}}_1} \phi, u - \hat{\mathbf{c}}_1 \rangle] \mathbb{1}_{A^c} \\ &\leq \mathbb{P}(A^c) P\phi(u) + \mathbb{E} \left[\sup_{c \in F_0 \cap \overline{\mathcal{B}}(0, \frac{k}{nh} \sum_{i=1}^n \|X_i\|)} (|\phi(c)| + P \|\nabla_c \phi\| (P \|u\| + \|c\|)) \right] \mathbb{1}_{A^c} \\ &\leq \mathbb{P}(A^c) P\phi(u) + \mathbb{E} \left[\left(2 \frac{k}{nh} \sum_{i=1}^n \|X_i\| + P \|u\| \right) \psi \left(\frac{k}{nh} \sum_{i=1}^n \|X_i\| \right) \right] \mathbb{1}_{A^c} \\ &\leq \mathbb{P}(A^c) P\phi(u) + \mathbb{E} \left[\frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \left(2 \frac{k \|X_i\|}{h} + P \|u\| \right) \psi \left(\frac{k \|X_j\|}{h} \right) \right] \mathbb{1}_{A^c} \\ &\leq \mathbb{P}(A^c) P\phi(u) + \sqrt{P(A^c)} \sqrt{C}, \end{aligned}$$

with

$$C = \max \left\{ \mathbb{E} \psi \left(\frac{k \|X\|}{h} \right)^2 \left(2(P\|u\|)^2 + \frac{8k^2}{h^2} \|X\|^2 \right), \mathbb{E} \psi \left(\frac{k \|X\|}{h} \right)^2 \mathbb{E} \left(2(P\|u\|)^2 + \frac{8k^2}{h^2} \|X\|^2 \right) \right\}$$

Since $\mathbb{P}(A^c) \leq \frac{1}{n}$, it follows that

$$\mathbb{E} \hat{\Delta} \mathbb{1}_{A^c} \leq \frac{C_P}{\sqrt{n}}.$$

Hence the result.

Proof of Proposition 3.19

Set $h \in (0, 1)$. From Lemma 3.12, we get that

$$d_{\phi, P, h}^2(\mathbf{c}') - d_{\phi, P, h}^2(\mathbf{c}) \leq \sum_{i=1}^k P_{\mathbf{c}, h}^{BV, i, \mathbf{c}}(\Omega) \omega_{\nabla \phi, \mathcal{K}}(\|c_i - c'_i\|) \left\| m \left(P_{\mathbf{c}, h}^{BV, i, \mathbf{c}} \right) - c'_i \right\|.$$

We now prove that $\left\| m \left(P_{\mathbf{c}, h}^{BV, i, \mathbf{c}} \right) \right\|$ is bounded when $\mathbf{c} \in \mathcal{K}^{(k)}$. Since P is inner-regular, we can set \mathcal{K}_P a convex compact subset of $\Omega \cap \mathcal{B}(0, K)$ for some $K > 0$ such that $P(\mathcal{K}_P) > h$. Take $\mathbf{c} \in \mathcal{K}^{(k)}$. Then, for every $y \in \mathcal{K}_P$ we get that:

$$\begin{aligned} d_{\phi}(y, \mathbf{c}) &= \min_{c \in \mathbf{c}} \phi(y) - \phi(c) - \langle \nabla_c \phi, y - c \rangle \\ &\leq \sup_{y \in \mathcal{K}_P, c \in \mathcal{K}} \|\phi(y) - \phi(c)\| + (K_P + K') \sup_{c \in \mathcal{K}} \|\nabla_c \phi\| \\ &\leq 2(K_P + K') \sup_{c \in \text{Conv}(\mathcal{K} \cup \mathcal{K}_P)} \|\nabla_c \phi\| \end{aligned}$$

according to the mean value theorem on the convex compact set $\text{Conv}(\mathcal{K} \cup \mathcal{K}_P)$. Note that $\sup_{c \in \text{Conv}(\mathcal{K} \cup \mathcal{K}_P)} \|\nabla_c \phi\| < +\infty$ since ϕ is C^1 . It holds that

$$\mathcal{K}_P \subset \mathcal{B}_{\phi} \left(\mathbf{c}, \sqrt{2(K_P + K') \sup_{c \in \text{Conv}(\mathcal{K} \cup \mathcal{K}_P)} \|\nabla_c \phi\|} \right).$$

Thus, $\delta_{\phi, P, h}(\mathbf{c}) \leq \Delta_{\phi, P, \mathcal{K}}$ with

$$\Delta_{\phi, P, \mathcal{K}} = \sqrt{2(K_P + K') \sup_{c \in \text{Conv}(\mathcal{K} \cup \mathcal{K}_P)} \|\nabla_c \phi\|},$$

which is finite.

Let $\epsilon > 0$ such that $\overline{\mathcal{B}(\mathcal{K}, \epsilon)} \subset \Omega$. It remains to prove that $\bigcup_{\mathbf{c} \in \mathcal{K}^{(k)}} \mathcal{B}_{\phi}(\mathbf{c}, \Delta_{\phi, P, \mathcal{K}})$ is included in $\mathcal{B}(0, \epsilon \Lambda_{\phi, P, \mathcal{K}, \epsilon})$, for some finite $\Lambda_{\phi, P, \mathcal{K}, \epsilon}$ to be determined.

Then, for all $v \in S(0, \epsilon)$, $c \in \mathcal{K}$ and $\lambda \geq 1$, we have:

$$d_{\phi}(c + \lambda v, c) \geq \phi(c + v) - \phi(c) + \langle \nabla_{c+v} \phi, (\lambda - 1)v \rangle - \langle \nabla_c \phi, \lambda v \rangle$$

since $d_{\phi}(c + \lambda v, c + v) \geq 0$. Thus, if $y = c + \lambda v \in \mathcal{B}_{\phi}(c, \Delta_{\phi, P, \mathcal{K}})$, then it holds that:

$$\lambda \leq \frac{\Delta_{\phi, P, \mathcal{K}}^2 + d_{\phi}(c, c + v)}{\langle \nabla_{c+v} \phi - \nabla_c \phi, v \rangle} \leq 1 + \frac{\Delta_{\phi, P, \mathcal{K}}^2}{\inf_{v \in S(0, \epsilon)} d_{\phi}(c, c + v) + d_{\phi}(c + v, c)}.$$

Set $\Lambda_{\phi, P, \mathcal{K}} = \left(1 + \frac{\Delta_{\phi, P, \mathcal{K}}^2}{\inf_{v \in S(0, \epsilon), c \in \mathcal{K}} d_{\phi}(c, c + v) + d_{\phi}(c + v, c)} \right)$. Then $\Lambda_{\phi, P, \mathcal{K}, \epsilon} < +\infty$ since $(c, v) \mapsto d_{\phi}(c, c + v) + d_{\phi}(c + v, c)$ is positive and continuous on the compact set $\mathcal{K} \times S(0, \epsilon)$, thus bounded from above by a positive constant, and $\bigcup_{\mathbf{c} \in \mathcal{K}^{(k)}} \mathcal{B}_{\phi}(\mathbf{c}, \Delta_{\phi, P, \mathcal{K}}) \subset \mathcal{B}(0, \epsilon \Lambda_{\phi, P, \mathcal{K}, \epsilon})$. Then, $L_{\phi, P, \mathcal{K}} = \epsilon \Lambda_{\phi, P, \mathcal{K}} + K'$ suits.

3.3.2 Proofs for Section 3.2

Proof of Proposition 3.26

When \mathcal{F} is a family of multivariate Gaussian distributions $P_{\theta, \Sigma}$ with covariance matrix $\Sigma = AA^T$, the distance to the measure $P = P_{\theta_0, \Sigma_0}$ with respect to the metric $\|\cdot\|_{\Sigma^{-1}}$ rewrites

$$d_{P, h, \|\cdot\|_{\Sigma^{-1}}}(\theta) = d_{\mathcal{N}(0, I_d), h, \|\cdot\|_{(A^{-1}A_0)^T(A^{-1}A_0)}}(A_0^{-1}(\theta - \theta_0)), \quad (3.22)$$

which is minimal at $\theta = \theta_0$. Equation (3.22) comes from the fact that $Y = A_0^{-1}(X - \theta_0) \sim \mathcal{N}(0, I_d)$ when $X \sim P_{\theta_0, \Sigma_0}$, and the fact that for all $x \in \mathbb{R}^d$, with $y = A_0^{-1}(x - \theta_0)$:

$$\begin{aligned} \|x - \theta\|_{\Sigma^{-1}} &\leq \delta_{P, h, \|\cdot\|_{\Sigma^{-1}}}(\theta) \\ &\iff \mathbb{P}_{X \sim P}(\|X - \theta\|_{\Sigma^{-1}} \leq \|x - \theta\|_{\Sigma^{-1}}) \leq h \\ &\iff \mathbb{P}_{Y \sim \mathcal{N}(0, I_d)}\left(\|Y + A_0^{-1}(\theta_0 - \theta)\|_{(A^{-1}A_0)^T(A^{-1}A_0)} \leq \|y + A_0^{-1}(\theta_0 - \theta)\|_{(A^{-1}A_0)^T(A^{-1}A_0)}\right) \leq h \\ &\iff \|y + A_0^{-1}(\theta_0 - \theta)\|_{(A^{-1}A_0)^T(A^{-1}A_0)} \leq \delta_{\mathcal{N}(0, 1), h, \|\cdot\|_{(A^{-1}A_0)^T(A^{-1}A_0)}}(A_0^{-1}(\theta - \theta_0)). \end{aligned}$$

Thus, $\delta_{P, h, \|\cdot\|_{\Sigma^{-1}}}(\theta) = \delta_{\mathcal{N}(0, 1), h, \|\cdot\|_{(A^{-1}A_0)^T(A^{-1}A_0)}}(A_0^{-1}(\theta - \theta_0))$ and (3.22) follows from Definition 2.1.

With the notation $\tilde{\Sigma} = ((A^{-1}A_0)^T(A^{-1}A_0))^{-1}$, it holds that $\det(\Sigma) = \det(\tilde{\Sigma}) \det(\Sigma_0)$, and the trimmed log-likelihood at (μ, Σ) expresses as

$$L_{P, h}((\mu, \Sigma)) = -\frac{d}{2} \log(2\pi) - \frac{1}{2} \log(\det(\tilde{\Sigma})) - \frac{1}{2} \log(\det(\Sigma_0)) - \frac{1}{2} d_{\mathcal{N}(0, I_d), h, \|\cdot\|_{\tilde{\Sigma}^{-1}}}^2(A_0^{-1}(\theta - \theta_0)),$$

and is maximal for $\theta = \theta_0$.

Now set $\Sigma' = \frac{\tilde{\Sigma}}{\det(\tilde{\Sigma})^{\frac{1}{d}}}$, then $\det(\Sigma') = 1$ and

$$d_{\mathcal{N}(0, I_d), h, \|\cdot\|_{\Sigma'^{-1}}}^2(0) = \frac{d_{\mathcal{N}(0, I_d), h, \|\cdot\|_{\tilde{\Sigma}^{-1}}}^2(0)}{\det(\tilde{\Sigma})^{\frac{1}{d}}}.$$

According to Lemma 3.27, $d_{\mathcal{N}(0, I_d), h, \|\cdot\|_{\Sigma'^{-1}}}^2(0)$ is minimal for $\Sigma' = I_d$. Thus, it remains to maximize

$$-\frac{d}{2} \log(2\pi) - \frac{1}{2} \log(\det(\tilde{\Sigma})) - \frac{1}{2} \log(\det(\Sigma_0)) - \frac{d_{\mathcal{N}(0, I_d), h, \|\cdot\|}^2(0)}{2 \det(\tilde{\Sigma})^{\frac{1}{d}}},$$

the maximum it attained for $\det(\tilde{\Sigma}) = \left(\frac{d_{\mathcal{N}(0, I_d), h, \|\cdot\|}^2(0)}{d}\right)^d$, which concludes.

Proof of Lemma 3.27

First recall that $d_{\mathcal{N}(0, 1), h, \|\cdot\|_{\Sigma^{-1}}}^2(0) = \frac{1}{h} \int_{l=0}^h \delta_{\mathcal{N}(0, 1), l, \|\cdot\|_{\Sigma^{-1}}}^2(0) dl$. In order to prove Lemma 3.27, it suffices to prove that for every $l > 0$, $\delta_{\mathcal{N}(0, 1), l, \|\cdot\|_{\Sigma^{-1}}}^2(0) \geq \delta_{\mathcal{N}(0, 1), l, \|\cdot\|}^2(0)$. Thus, it suffices to prove that:

$$\forall \delta > 0, \mathbb{P}\left(\sum_{i=1}^n \lambda_i X_i^2 < \delta\right) \leq \mathbb{P}\left(\sum_{i=1}^n X_i^2 < \delta\right) \quad (3.23)$$

when $\prod_{i=1}^n \lambda_i = 1$ and the λ_i s are positive, where the probability \mathbb{P} is computed according to X_1, X_2, \dots, X_n iid from $\mathcal{N}(0, 1)$. Indeed, the matrix Σ^{-1} can be diagonalised in an orthogonal basis, and the λ_i s correspond to its eigenvalues. Equation 3.23 can be proved as follows.

Set $\mathcal{B}_\lambda(\delta) = \{x_1, x_2, \dots, x_n \in \mathbb{R} \mid \sum_{i=1}^n \lambda_i x_i^2 < \delta\}$, and $\mathcal{B}(\delta) = \{x_1, x_2, \dots, x_n \in \mathbb{R} \mid \sum_{i=1}^n x_i^2 < \delta\}$. Since $\prod_{i=1}^n \lambda_i = 1$, after an integration by substitution, it follows that $\text{Leb}(\mathcal{B}_\lambda(\delta)) = \text{Leb}(\mathcal{B}(\delta))$, where Leb denotes the Lebesgue measure on \mathbb{R}^d . Then, for all $\delta > 0$, if $f(\delta)$ denotes the density of $P = \mathcal{N}(0, I_d)$ at any point $x = (x_1, x_2, \dots, x_n)$ such that $\sum_{i=1}^n x_i^2 = \delta$, it holds:

$$\begin{aligned} P(\mathcal{B}_\lambda(\delta) \setminus \mathcal{B}(\delta)) &\geq f(\delta) \text{Leb}(\mathcal{B}_\lambda(\delta) \setminus \mathcal{B}(\delta)) \\ &= f(\delta) \text{Leb}(\mathcal{B}(\delta) \setminus \mathcal{B}_\lambda(\delta)) \\ &\geq P(\mathcal{B}(\delta) \setminus \mathcal{B}_\lambda(\delta)), \end{aligned}$$

which concludes.

Proof of Proposition 3.29

The family of all distributions $\mathcal{N}(\mu, \sigma^2)$ for $\mu \in \mathbb{R}$ and $\sigma > 0$ can be expressed as an exponential family with parameters $\theta = (\frac{\mu}{\sigma^2}, \frac{-1}{2\sigma^2}) = \theta(\mu, \sigma^2) \in \mathbb{R}^2$, $f(x) = (x, x^2) \in \mathbb{R}^2$ and $\psi(\theta) = -\frac{\theta_1^2}{4\theta_2} + \frac{1}{2} \log(-\frac{\pi}{\theta_2})$, that is, $\psi(\theta(\mu, \sigma^2)) = \frac{1}{2} \left(\frac{\mu^2}{\sigma^2} + \log(2\pi) + \log(\sigma^2) \right)$. Indeed, $p_{\theta(x)} = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right)$.

Then, for $\theta = \theta(\mu, \sigma^2)$, $H(\theta, c_{P,h}(\theta)) = \mathcal{B}(\mu, \delta_{P,h}(\mu))$, with

$$\delta_{P,h}(\mu) = \inf \{r > 0 \mid P(\overline{\mathcal{B}}(\mu, r)) > h\}.$$

It does not depend on $\chi = \sigma^2$. It means that for all $\chi \in \Xi$, $\mathcal{P}_{P\theta(u,\chi),h}(P) = \mathcal{P}_{\mu_0,h}(P)$, when $P = P_{\theta_0}$, we get that $\theta^* = P\theta(u, \chi^*) = \theta(\mu_0, \chi^*)$. It remains to compute χ^* . We may write,

$$\begin{aligned} &P(P_{\mu_0,h} \langle \theta(v, \chi), f(u) \rangle - \psi(\theta(v, \chi))) du dv \\ &= PP_{\mu_0,h} \left\langle \left(\frac{v}{\sigma^2}, \frac{-1}{2\sigma^2} \right), (u, u^2) \right\rangle dudv - \frac{1}{2} \frac{v^2}{\sigma^2} - \frac{1}{2} \log(2\pi) - \frac{1}{2} \log(\sigma^2) \\ &= PP_{\mu_0,h} - \frac{(u-v)^2}{2\sigma^2} - \frac{1}{2} \log(2\pi) - \frac{1}{2} \log(\sigma^2) \\ &= \frac{P_{\mu_0,h} - ((u-\mu_0)^2 + \sigma_0^2) du}{2\sigma^2} - \frac{1}{2} \log(2\pi) - \frac{1}{2} \log(\sigma^2) \\ &= -\frac{1}{2} \log(2\pi) - \frac{1}{2} \log \sigma^2 - \frac{d_{P,h}^2(\mu_0) + \sigma_0^2}{2\sigma^2}. \end{aligned}$$

Since $d_{P,h}(\mu_0) = \sigma_0 d_{\mathcal{N}(0,1),h}(0)$ it comes that the maximum is attained at $\chi^* = \sigma^{*2} = (1 + d_{\mathcal{N}(0,1),h}^2(0)) \sigma_0^2$.

Proof of Proposition 3.30

The family of all distributions $\mathcal{N}(\mu, \Sigma = AA^T)$ for $\mu \in \mathbb{R}^d$ and A an invertible $d \times d$ matrix with elements in \mathbb{R} , can be expressed as an exponential family with parameters $\theta = (\Sigma^{-1}\mu, \text{vect}(-\frac{1}{2}\Sigma^{-1})) = \theta(\mu, A)$, $f(x) = (x, \text{vect}(xx^T))$ and $\psi(\theta) = \frac{1}{2}\mu^T \Sigma^{-1}\mu + \frac{1}{2} \log(\det(2\pi\Sigma))$. Let $P = P_{\mu_0, \Sigma_0}$, then, $P\theta(u, \Sigma) = \theta(\mu_0, \Sigma)$, moreover,

$$\begin{aligned} &P \langle \theta(v, \Sigma), P_{\theta(\mu_0, \Sigma), h} f(u) \rangle - \psi(\theta(v, \Sigma)) dv \\ &= -PP_{\theta(\mu_0, \Sigma), h} \frac{1}{2} (u-v)^T \Sigma^{-1} (u-v) - \frac{1}{2} \log(\det(2\pi\Sigma)) dudv \\ &= -P_{\theta(\mu_0, \Sigma), h} \left(P \frac{1}{2} (\mu_0 - v)^T \Sigma^{-1} (\mu_0 - v) dv \right) - \frac{1}{2} (u - \mu_0)^T \Sigma^{-1} (u - \mu_0) du - \frac{1}{2} \log(\det(2\pi\Sigma)) \\ &= -\frac{1}{2} \text{Tr}(\Sigma^{-1}\Sigma_0) - \frac{1}{2} P_{\theta(\mu_0, \Sigma), h} (u - \mu_0)^T \Sigma^{-1} (u - \mu_0) du - \frac{1}{2} \log(\det(2\pi\Sigma)) \end{aligned}$$

where

$$\begin{aligned} P_{\theta(\mu_0, \Sigma), h}(u - \mu_0)^T \Sigma^{-1}(u - \mu_0) du &= d_{P, h, \|\cdot\|_{\Sigma^{-1}}}^2(\mu_0) \\ &= d_{\mathcal{N}(0, I_d), h, \|\cdot\|_{(A^{-1}A_0)^T(A^{-1}A_0)}}^2(0), \end{aligned}$$

according to (3.22). Set $\tilde{\Sigma} = ((A^{-1}A_0)^T(A^{-1}A_0))^{-1}$, and $\Sigma' = \frac{\tilde{\Sigma}}{\det(\tilde{\Sigma})^{\frac{1}{d}}}$.

Then it remains to minimise

$$\begin{aligned} &\frac{Tr(\Sigma'^{-1})}{\det(\tilde{\Sigma})^{\frac{1}{d}}} + d_{\mathcal{N}(0, I_d), h, \|\cdot\|_{\tilde{\Sigma}^{-1}}}^2(0) + \log(\det(\Sigma)) \\ &= \frac{Tr(\Sigma'^{-1})}{\det(\tilde{\Sigma})^{\frac{1}{d}}} + \frac{d_{\mathcal{N}(0, I_d), h, \|\cdot\|_{\Sigma'^{-1}}}^2(0)}{\det(\tilde{\Sigma})^{\frac{1}{d}}} + d \log(\det(\tilde{\Sigma})^{\frac{1}{d}}) + \log(\det(\Sigma_0)). \end{aligned}$$

Then the minimum is attained for $\det(\tilde{\Sigma})^{\frac{1}{d}} = \frac{Tr(\Sigma'^{-1}) + d_{\mathcal{N}(0, I_d), h, \|\cdot\|_{\Sigma'^{-1}}}^2(0)}{d}$ and according to Lemma 3.27, $d_{\mathcal{N}(0, I_d), h, \|\cdot\|_{\Sigma'^{-1}}}^2(0)$ is minimal at $\Sigma' = I_d$.

Moreover $Tr(\Sigma'^{-1})$ is also minimal for $\Sigma' = I_d$, according to the inequality of arithmetic and geometric means for the eigenvalues of Σ' , which product equals to 1. Thus the minimum is unique and satisfies $\Sigma^* = \frac{d + d_{\mathcal{N}(0, I_d), h, \|\cdot\|}^2(0)}{d} \Sigma_0$.

Proof of Proposition 3.31

For every $\mu = (\mu_1, \mu_2, \dots, \mu_k) \in \mathcal{X}^{(k)}$ and $\chi = (\chi_1, \chi_2, \dots, \chi_k) \in \Xi^{(k)}$, we note $\theta_i = \theta(\mu_i, \chi_i)$,

$$\begin{aligned} \theta'_i &= \frac{1}{|\mathcal{C}(\theta_i, \chi_i)|} \sum_{X \in \mathcal{C}(\theta_i, \chi_i)} \theta(X, \chi_i), \\ \chi'_i &\in \arg \max \left\{ \frac{1}{|\mathcal{C}(\theta_i, \chi_i)|} \sum_{X \in \mathcal{C}(\theta_i, \chi_i)} \langle \theta(X, \chi), P_{n\theta'_i, h} f(u) \rangle - \psi(\theta(X, \chi)) \mid \chi \in \mathcal{X} \right\}, \\ \theta''_i &= \frac{1}{|\mathcal{C}(\theta_i, \chi_i)|} \sum_{X \in \mathcal{C}(\theta_i, \chi_i)} \theta(X, \chi'_i), \end{aligned}$$

$\mathcal{C}_i = \mathcal{C}(\theta_i, \chi_i)$ and $\mathcal{C}'_i = \mathcal{C}(\theta''_i, \chi'_i)$. Then,

$$\begin{aligned}
& \frac{1}{n} \sum_{X \in \mathbb{X}_n} \min_{i \in [1, k]} \langle \theta(X, \chi_i), P_n \theta_i, h f(u) \rangle - \psi(\theta(X, \chi_i)) \\
&= \frac{1}{n} \sum_{i=1}^k \left(\sum_{X \in \mathcal{C}_i} (-\psi(\theta(X, \chi_i)) + |\mathcal{C}_i| \langle \theta'_i, P_n \theta_i, h f(u) \rangle) \right) \\
&\leq \frac{1}{n} \sum_{i=1}^k \sum_{X \in \mathcal{C}_i} \left(-\psi(\theta(X, \chi_i)) + \langle \theta(X, \chi_i), P_n \theta'_i, h f(u) \rangle \right) \\
&\leq \frac{1}{n} \sum_{i=1}^k \sum_{X \in \mathcal{C}_i} \left(-\psi(\theta(X, \chi'_i)) + \langle \theta(X, \chi'_i), P_n \theta'_i, h f(u) \rangle \right) \\
&= \frac{1}{n} \sum_{i=1}^k \left(\sum_{X \in \mathcal{C}_i} (-\psi(\theta(X, \chi'_i)) + |\mathcal{C}_i| \langle \theta''_i, P_n \theta'_i, h f(u) \rangle) \right) \\
&\leq \frac{1}{n} \sum_{i=1}^k \sum_{X \in \mathcal{C}_i} \left(-\psi(\theta(X, \chi'_i)) + \langle \theta(X, \chi'_i), P_n \theta''_i, h f(u) \rangle \right) \\
&\leq \frac{1}{n} \sum_{i=1}^k \sum_{X \in \mathcal{C}'_i} \left(-\psi(\theta(X, \chi'_i)) + \langle \theta(X, \chi'_i), P_n \theta''_i, h f(u) \rangle \right) \\
&= \frac{1}{n} \sum_{X \in \mathbb{X}_n} \max_{i \in [1, k]} \langle \theta(X, \chi'_i), P_n \theta''_i, h f(u) \rangle - \psi(\theta(X, \chi'_i))
\end{aligned}$$

We used twice (3.11) to get for instance that $\langle \theta'_i, P_n \theta'_i, h f(u) \rangle \geq \langle \theta'_i, P_n \theta_i, h f(u) \rangle$.

Proof of Lemma 3.33

$$\begin{aligned}
\psi(\Sigma) &= QQ' \langle \theta(u, \Sigma), f(v) \rangle - \psi(\theta(u, \Sigma)) dudv \\
&= QQ' \frac{1}{2} \left(-(v-u)^T \Sigma^{-1} (v-u) - \log(\det(2\pi\Sigma)) \right) dudv
\end{aligned}$$

The matrix Σ is symmetric with coefficients in \mathbb{R} , thus diagonalisable in an orthogonal basis. Thus, we can write

$$\Sigma = PDP^{-1} = PDP^T$$

for some diagonal matrix $D = (d_{i,j})_{i,j}$ with $d_{i,j} = \delta_{i,j} \lambda_i$. Here $\delta_{i,j} = 1$ when $i = j$ and 0 when $i \neq j$. We note $P = [p_{i,j}]_{i,j}$. Then,

$$\begin{aligned}
-\psi(\Sigma) &= QQ' \frac{1}{2} \left((v-u)^T PD^{-1}P^T(v-u) + \log((2\pi)^d \det(PDP^{-1})) \right) \\
&= \frac{1}{2} QQ' \left(\log(2\pi)^d + \sum_{i,j=1}^d (v_i - u_i)(v_j - u_j) \left(\sum_{k=1}^d p_{j,k} p_{i,k} \lambda_k^{-1} \right) + \sum_{k=1}^d \log(\lambda_k) \right) \\
&= \frac{1}{2} \left(\log(2\pi)^d + \sum_{k=1}^d \lambda_k^{-1} \left[QQ' \sum_{i,j=1}^d (v_i - u_i)(v_j - u_j) p_{j,k} p_{i,k} \right] + \log \lambda_k \right)
\end{aligned}$$

Thus, $\psi(\Sigma)$ is maximal when $\lambda_k = \lambda_k^*$, where:

$$\lambda_k^* = QQ' \sum_{i,j=1}^d (v_i - u_i)(v_j - u_j) p_{j,k} p_{i,k} = [P^T AP]_{k,k},$$

where A is defined by:

$$A = [QQ' \sum_{i,j=1}^d (v_i - u_i)(v_j - u_j)]_{i,j}.$$

Thus, it remains to minimise the function

$$\tilde{\psi} : P \mapsto Tr(\log(P^T AP)) = \sum_{k=1}^d \log \left(\sum_{i,j=1}^d [QQ'(v_i - u_i)(v_j - u_j)] p_{j,k} p_{i,k} \right).$$

Again, we can diagonalise A in an orthonormal basis: $A = P_0 D_0 P_0^T = P_0 D_0 P_0^{-1}$. By setting $\tilde{P} = P^T P_0$, we get

$$\begin{aligned} \tilde{\psi}(P_0 \tilde{P}^{-1}) &= Tr \log(\tilde{P} D_0 \tilde{P}^T) \\ &= \sum_{l=1}^d \log \left(\sum_{k=1}^d \tilde{p}_{l,k}^2 \lambda_k^0 \right) \\ &\geq \sum_{l=1}^d \sum_{k=1}^d \tilde{p}_{l,k}^2 \log(\lambda_k^0) \\ &= \sum_{k=1}^d \log(\lambda_k^0) \sum_{l=1}^d \tilde{p}_{l,k}^2 \\ &= Tr \log[I_d D_0 I_d^T]. \end{aligned}$$

We used concavity of the function \log , and the fact that \tilde{P} is orthogonal, thus, $\sum_{l=1}^d \tilde{p}_{l,k}^2 = \sum_{k=1}^d \tilde{p}_{l,k}^2 = 1$. Thus, we can choose $\tilde{P} = I_d$, that is, $P = P_0$. Moreover, $\lambda_k^* = [P_0^T A P_0]_{k,k} = \lambda_k^0$. Thus, $\Sigma = P_0 D_0 P_0^T = A$ is a minimizer of ψ .

Appendix

3.A Maximal deviation bounds for empirical processes – Background and first results

Dans ce Chapitre, on présente deux méthodes bien connues en statistiques pour calculer des bornes supérieures pour les maxima de déviation de processus empiriques de type $\sup_{f \in \mathcal{F}} |P_n - P|f(u)$, lorsque P est une mesure Borélienne de probabilité sur un espace métrique Polonais \mathcal{X} et \mathcal{F} un ensemble de fonctions mesurables bornées sur \mathcal{X} à valeurs dans \mathbb{R} . On propose des exemples d'application de ces méthodes qui seront utilisés dans les chapitres suivants pour obtenir des vitesses de convergence des estimateurs de codebooks optimaux. Ces méthodes reposent sur le théorème des différences bornées. Afin de borner l'espérance du maximum de déviation, on utilisera des arguments de symétrisation faisant entrer en jeu des moyennes de Rademacher. Enfin, on utilisera des notions de dimension de Vapnik-Chervonenkis d'une part et de nombre de recouvrement, d'entropie métrique avec le théorème de l'intégrale de l'entropie de Dudley d'autre part, selon que \mathcal{F} soit un ensemble d'indicatrices ou un ensemble de fonctions bornées. La seconde méthode est aussi connue sous le nom de chainage. Let P be a Borel probability distribution on a Polish metric space (\mathcal{X}, δ) . We call n -sample from P any n -tuple $\mathbb{X}_n = (X_1, X_2, \dots, X_n)$ of independent random variables X_1, X_2, \dots, X_n sampled according to P . The distribution of such an n -sample is denoted by \mathbb{P} , and the expectation of a function f of (X_1, X_2, \dots, X_n) by \mathbb{E} . Also, $P_n = \frac{1}{n} \sum_{i=1}^n \delta_{X_i}$ denotes the *empirical distribution* associated to the n -sample \mathbb{X}_n from P . As Pf or $Pf(u)$ denotes the expectation of $f(X)$ when X is sampled from P , we will use the notation $P_n f = P_n f(u) = \frac{1}{n} \sum_{i=1}^n f(X_i)$ for the expectation of f for the measure P_n .

Given a countable family of measurable functions \mathcal{F} , [BBL05, Section 3] and [BLM13] are main references to derive upper bounds for the maximal deviation

$$\sup_{f \in \mathcal{F}} |Pf - P_n f|,$$

with the use of symmetrisation arguments, Rademacher averages and Vapnik dimension, or also metric entropy and Dudley integral, with many references therein.

In this chapter, we expose two methods to derive upper bounds for the maximal deviations, with high probability. We also provide examples of applications of these methods. Such

examples will be used as lemmas to derive rates of convergence of estimators of the optimal codebooks in the following chapters. The first method is adapted to the particular case when \mathcal{F} is a set of indicators. The bounds obtained depend on the Vapnik-Chervonenkis dimension. The second method is adapted to all sets \mathcal{F} of bounded functions and relies on the Dudley entropy integral. The bounds obtained depend on metric entropy and covering numbers. This second method is also known as chaining. Both methods start with an application of the bounded difference inequality. Then, symmetrisation arguments with Rademacher averages are used to derive upper bounds for the expectation $\mathbb{E} [\sup_{f \in \mathcal{F}} |Pf - P_n f|]$.

3.A.1 From bounded difference inequality to Rademacher averages

Let \mathcal{X} a measurable space. A function $g : \mathcal{X}^n \mapsto \mathbb{R}$ has the *bounded differences property* if for some non-negative constants c_1, c_2, \dots, c_n , for all $i \in \llbracket 1, n \rrbracket$:

$$\sup_{x_1, x_2, \dots, x_n, x'_i \in \mathcal{X}} |g(x_1, \dots, x_i, \dots, x_n) - g(x_1, \dots, x'_i, \dots, x_n)| \leq c_i.$$

Theorem 3.43 (Bounded difference inequality [BLM13, Theorem 6.2] and [BBL05, Theorem 3.1]). *Assume that the function g satisfies the bounded differences assumption with constants c_1, \dots, c_n and denote*

$$v = \frac{1}{4} \sum_{i=1}^n c_i^2$$

Let $Z = g(X_1, \dots, X_n)$ where the X_i are independent. Then

$$P\{Z - EZ > t\} \leq \exp\left(-\frac{t^2}{2v}\right).$$

As a consequence, bounds are derived for a random variable $Z = \sup_{f \in \mathcal{F}} |Pf - P_n f|$ when f are non-negative bounded measurable functions and \mathcal{F} a countable set.

Corollary 3.44. *Set \mathcal{F} a countable family of measurable functions $f : \mathcal{X} \rightarrow \mathbb{R}_+$ bounded by $R > 0$ from above. Set $Z = \sup_{f \in \mathcal{F}} |Pf - P_n f|$. Then,*

$$P\{Z - \mathbb{E}[Z] > t\} \leq \exp\left(-\frac{2nt^2}{R^2}\right).$$

Proof. The function $g : (x_1, x_2, \dots, x_n) \mapsto \sup_{f \in \mathcal{F}} |Pf(u) - \frac{1}{n} \sum_{i=1}^n f(x_i)|$ satisfies the bounded difference inequality with parameter $c_i = \frac{R}{n}$. The inequality follows from Theorem 3.43. \square

Note that when f takes values in $[-R, R]$, the upper-bound is in fact $\exp\left(-\frac{nt^2}{2R^2}\right)$. This corollary can be extended to separable metric spaces \mathcal{F} , with \mathcal{F}' a dense countable subset of \mathcal{F} provided that

$$\sup_{f \in \mathcal{F}} |Pf - P_n f| = \sup_{f \in \mathcal{F}'} |Pf - P_n f| \text{ a.e.}$$

As a consequence, provided that \mathcal{F} is a countable family of non-negative functions bounded from above by $R > 0$, for every $\delta > 0$, the following inequality holds with probability at least $1 - \delta$:

$$\sup_{f \in \mathcal{F}} |Pf - P_n f| \leq \mathbb{E} \left[\sup_{f \in \mathcal{F}} |Pf - P_n f| \right] + \sqrt{\frac{R^2}{2n} \log \frac{1}{\delta}}. \quad (3.24)$$

When \mathcal{F} is a countable family of functions taking values in $[-R, R]$, the following inequality holds with probability at least $1 - \delta$:

$$\sup_{f \in \mathcal{F}} |Pf - P_n f| \leq \mathbb{E} \left[\sup_{f \in \mathcal{F}} |Pf - P_n f| \right] + \sqrt{\frac{2R^2}{n} \log \frac{1}{\delta}}. \quad (3.25)$$

It remains to derive upper-bounds for the quantity $\mathbb{E} [\sup_{f \in \mathcal{F}} |Pf - P_n f|]$. For this we use a symmetrization argument with the help of Rademacher independent random variables $\sigma_1, \sigma_2 \dots \sigma_n$, independent from the X_i s. The *Rademacher average* associated to a countable set A with elements in \mathbb{R}^d is defined by:

$$R_n(A) = \mathbb{E}_\sigma \sup_{a \in A} \frac{1}{n} \left| \sum_{i=1}^n \sigma_i a_i \right|. \quad (3.26)$$

We define the subset $\mathcal{F}(x_1^n)$ of \mathbb{R}^n for all $x_1^n = (x_1, x_2, \dots, x_n) \in \mathcal{X}^n$ by

$$\mathcal{F}(x_1^n) = \{(f(x_1), f(x_2), \dots, f(x_n)) \mid f \in \mathcal{F}\}.$$

Then, for all countable family \mathcal{F} of bounded measurable functions, it holds (see e.g. [BBL05, Theorem 3.2])

$$\mathbb{E} \left[\sup_{f \in \mathcal{F}} |Pf - P_n f| \right] \leq 2\mathbb{E} [R_n(\mathcal{F}(X_1^n))]. \quad (3.27)$$

When \mathcal{F} is finite, it comes from the Jensen inequality applied to the convex function $x \mapsto \sup_{i \in [1, n]} x_i$ on \mathbb{R}^n . The case where \mathcal{F} is a countable family of bounded measurable functions is a consequence of the case \mathcal{F} finite and the dominated convergence theorem.

It is also possible to define a slight modification of the Rademacher average, with no absolute value:

$$\tilde{R}_n(\mathcal{F}(X_1^n)) = \mathbb{E}_\sigma \sup_{f \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^n \sigma_i f(X_i). \quad (3.28)$$

Again, as a consequence of the Jensen inequality, see [BLM13, Lemma 11.4], it comes for all countable family of bounded measurable functions that:

$$\mathbb{E} \left[\sup_{f \in \mathcal{F}} Pf - P_n f \right] \leq 2\mathbb{E} [\tilde{R}_n(\mathcal{F}(X_1^n))]$$

Denote by $-\mathcal{F}$ the set of functions $-f$ when $f \in \mathcal{F}$, then we also have:

$$\mathbb{E} \left[\sup_{f \in \mathcal{F}} |Pf - P_n f| \right] \leq 2\mathbb{E} [\tilde{R}_n((\mathcal{F} \cup -\mathcal{F})(X_1^n))]. \quad (3.29)$$

In the following we first derive an upper bound for the Rademacher average $R_n(\mathcal{F}(X_1^n))$ when \mathcal{F} is a family of indicators, with the help of combinatorial tools such as the Vapnik Chervonenkis dimension. Then we derive an upper bound for the modified Rademacher average $\tilde{R}_n(\mathcal{F}(X_1^n))$ when \mathcal{F} is a set of bounded measurable functions. To this aim, we will introduce the metric entropy. As a consequence, we can derive upper-bounds for the probability of events such that $\sup_{x, r} |P(\mathcal{B}(x, r)) - P_n(\mathcal{B}(x, r))| \geq \epsilon$, $\sup_{x, r} \|Pu \mathbb{1}_{\mathcal{B}(x, r)}(u) - P_n u \mathbb{1}_{\mathcal{B}(x, r)}(u)\| \geq \epsilon$ or even $\sup_{x, r} |P\|u\|^2 \mathbb{1}_{\mathcal{B}(x, r)}(u) - P_n \|u\|^2 \mathbb{1}_{\mathcal{B}(x, r)}(u)| \geq \epsilon$ under assumptions on P . We also derive bounds for maximal deviations that involve Bregman balls and Bregman divergences.

3.A.2 \mathcal{F} , a family of indicators

The Vapnik-Chervonenkis inequality

First, we focus on properties for Rademacher averages defined by (3.26). From [BBL05, Theorem 3.3] it holds for all finite subset $A = \{a_1, a_2, \dots, a_{|A|}\}$ of cardinal $|A|$ of \mathbb{R}^d :

$$R_n(A) \leq \max_{i \in \llbracket 1, |A| \rrbracket} \|a_i\| \frac{\sqrt{2 \log(|A|)}}{n}. \quad (3.30)$$

In particular, when \mathcal{F} is a family of indicators, any element a of the set $A = \mathcal{F}(x_1^n)$ satisfies $\|a\| \leq \sqrt{n}$. Set $\mathbb{S}_{\mathcal{F}}(x_1^n)$ the cardinal of $\mathcal{F}(x_1^n)$ for $x_1^n = (x_1, x_2, \dots, x_n) \in \mathcal{X}^n$, that is called *VC shatter coefficient* (where VC stands for Vapnik-Chervonenkis), then it holds:

$$R_n(\mathcal{F}(x_1^n)) \leq \sqrt{\frac{2 \log(\mathbb{S}_{\mathcal{F}}(x_1^n))}{n}}. \quad (3.31)$$

Definition 3.45. The *VC dimension* of a subset A of $\{0, 1\}^n$, denoted by $V(A)$ is defined as the size V of the largest set of indices $\{i_1, i_2, \dots, i_V\} \subset \llbracket 1, n \rrbracket$ such that for each element $b \in \{0, 1\}^V$, there exists $a \in A$ such that $(a_{i_1}, a_{i_2}, \dots, a_{i_V}) = b$.

The Sauer's lemma states that $|A| \leq (n + 1)^V$. In particular,

$$\log(\mathbb{S}_{\mathcal{F}}(x_1^n)) \leq V(\mathcal{F}(x_1^n)) \log(n + 1). \quad (3.32)$$

The consequence of (3.31), (3.32) and (3.27) is the first inequality of the following lemme (see e.g. [BBL05, Theorem 3.4]).

Theorem 3.46 (Vapnik-Chervonenkis inequality). *For all countable family \mathcal{F} of indicator functions, if $V = \sup_{n, x_1^n} V(\mathcal{F}(x_1^n))$ where the supremum is taken on all possible values x_1^n taken by X_1^n , it holds:*

$$\mathbb{E} \left[\sup_{f \in \mathcal{F}} |Pf - P_n f| \right] \leq 2 \sqrt{\frac{2V \log(n + 1)}{n}}$$

and for some universal constant C :

$$\mathbb{E} \left[\sup_{f \in \mathcal{F}} |Pf - P_n f| \right] \leq C \sqrt{\frac{V}{n}}.$$

Proof. The first inequality is a consequence of (3.27), (3.31) and (3.32). The second inequality relies on chaining arguments is a refinement of the first one, see Theorem 3.55, (3.45), (3.39) and (3.44). \square

Computation of Vapnik-Chervonenkis dimension – some examples

Note that the Vapnik-Chervonenkis dimension of $\mathcal{F}(x_1^n)$ where \mathcal{F} is the set of all indicators or balls or of half-spaces in \mathbb{R}^d at most $d + 1$. Indeed, according to [Dud79], for any set E of $d + 2$ points in \mathbb{R}^d , not all subsets of E can be written as $E \cap B$ where B is a Euclidean ball. The same holds for half spaces.

It is also possible to generalise to Bregman-balls, whose Vapnik-Chervonenkis dimension is at most $d + 1$, as depicted below.

Proposition 3.47. *Let ϕ be a differentiable convex function. Set \mathcal{F} , the set of all indicators of Bregman-balls $\mathcal{B}_{\phi}(x, r) = \{y \in \mathbb{R}^d \mid \sqrt{d_{\phi}(y, x)} < r\}$, for $x \in \mathbb{R}^d$ and $r \geq 0$. Then, the Vapnik-Chervonenkis dimension of $\mathcal{F}(x_1^n)$ is at most $d + 1$ for all $x_1^n = (x_1, x_2, \dots, x_n) \in (\mathbb{R}^d)^n$.*

Proof. Let $S = \{x_1, \dots, x_{d+2}\}$ be shattered by \mathcal{F} . And let A_1, A_2 be a partition of S . Then for some $c_1 \in \mathbb{R}^d$ and $r_1 \geq 0$, it holds that $A_1 \subset \mathcal{B}_\phi(c_1, r_1)$ and $A_2 \subset \mathcal{B}_\phi^c(c_1, r_1)$. By exchanging the roles of A_1 and A_2 , we may write for some $c_2 \in \mathbb{R}^d$ and $r_2 \geq 0$

$$\begin{aligned} A_1 &= S \cap \mathcal{B}_\phi(c_1, r_1) \cap \mathcal{B}_\phi(c_2, r_2)^c \\ A_2 &= S \cap \mathcal{B}_\phi(c_2, r_2) \cap \mathcal{B}_\phi(c_1, r_1)^c. \end{aligned}$$

In particular, it holds for instance that for all $y \in A_1$, $d_\phi(y, c_i) = \phi(y) - \phi(c_i) - \langle \nabla_{c_i} \phi, y - c_i \rangle \leq r_1^2$. Straightforward computation shows that, for any $x \in A_1$,

$$\ell_{1,2}(x) < 0,$$

where $\ell_{1,2}(x) = \phi(c_2) - \phi(c_1) + \langle x, \nabla_{c_2} \phi - \nabla_{c_1} \phi \rangle + \langle \nabla_{c_1} \phi, c_1 \rangle - \langle \nabla_{c_2} \phi, c_2 \rangle - r_1^2 + r_2^2$. Similarly we have that, for any $x \in A_2$ $\ell_{1,2}(x) > 0$. Thus S is shattered by affine hyperplanes, whose VC-dimension is $d + 1$ according to [Dud79], hence the contradiction. \square

Application to \mathcal{F} a set of indicator of balls

As a direct consequence of the bounded inequality theorem together with the Vapnik-Chervonenkis inequality, it is possible to derive bounds for empirical processes of indicators of balls in $\frac{\log(n+1)}{\sqrt{n}}$:

Corollary 3.48. *Let $P \in \mathcal{P}(\mathbb{R}^d)$, then for all $p > 0$, with probability larger than $1 - 2n^{-p}$,*

$$\begin{aligned} \sup_{x \in \mathbb{R}^d, r > 0} |P(\overline{\mathcal{B}}(x, r)) - P_n(\overline{\mathcal{B}}(x, r))| &\leq 2\sqrt{\frac{2(d+1)\log(n+1)}{n}} + \sqrt{\frac{p \log n}{2n}}, \\ \sup_{x \in \mathbb{R}^d, r > 0} |P(\mathcal{B}(x, r)) - P_n(\mathcal{B}(x, r))| &\leq 2\sqrt{\frac{2(d+1)\log(n+1)}{n}} + \sqrt{\frac{p \log n}{2n}}, \end{aligned}$$

in particular,

$$\sup_{x \in \mathbb{R}^d, r > 0} |P(\partial \mathcal{B}(x, r)) - P_n(\partial \mathcal{B}(x, r))| \leq 4\sqrt{\frac{2(d+1)\log(n+1)}{n}} + 2\sqrt{\frac{p \log n}{2n}}.$$

The same holds for half-spaces.

Note that the term $2\sqrt{\frac{2(d+1)\log(n+1)}{n}}$ can be replaced everywhere by $C\sqrt{\frac{d+1}{n}}$ for some absolute positive constant C .

Proof. First set \mathcal{D} a dense and countable subset of the separable metric space $\mathbb{R}^d \times \mathbb{R}_+$. Since for all $x \in \mathbb{R}^d, r > 0$, $\overline{\mathcal{B}}(x, r)$ writes as a decreasing intersection of balls $\overline{\mathcal{B}}(x_n, r_n)$ with $(x_n, r_n) \in \mathcal{D}$,

$$\sup_{(x,r) \in \mathcal{D}} |P(\overline{\mathcal{B}}(x, r)) - P_n(\overline{\mathcal{B}}(x, r))| = \sup_{(x,r) \in \mathbb{R}^d \times \mathbb{R}_+} |P(\overline{\mathcal{B}}(x, r)) - P_n(\overline{\mathcal{B}}(x, r))|.$$

The same holds for the open balls. Thus there is no problem of measurability.

The inequalities follow from (3.24), Theorem 3.46 with $\delta = n^{-p}$ and [Dud79] for the Vapnik-Chervonenkis dimension of indicators of balls. \square

3.A.3 \mathcal{F} , a family of bounded functions

In this section, we present the chaining method.

Covering and packing numbers

Let (\mathcal{T}, d) be a totally bounded pseudo-metric space, meaning that for every $\delta > 0$, \mathcal{T} can be covered by a finite number of d -balls of radius δ . Given some $\delta > 0$, a δ -net denoted by \mathcal{T}_δ is a subset of \mathcal{T} with maximal cardinality such that for all $s, t \in \mathcal{T}_\delta$ with $s \neq t$, we have $d(s, t) > \delta$. In particular, it holds that:

$$\mathcal{T} \subset \bigcup_{t \in \mathcal{T}_\delta} \overline{\mathcal{B}}(t, \delta). \quad (3.33)$$

We denote by $N_d(\delta, \mathcal{T})$ the cardinality of a δ -net \mathcal{T}_δ and call it the δ -packing number of \mathcal{T} .

A notion related to the δ -packing number is the δ -covering number, it corresponds to the minimal cardinality of a set of points \mathcal{T}'_δ such that $\mathcal{T} \subset \bigcup_{t \in \mathcal{T}'_\delta} \overline{\mathcal{B}}(t, \delta)$. As noted in [BLM13], these notions are related as follows:

$$N_d(2\delta, \mathcal{T}) \leq N'_d(\delta, \mathcal{T}) \leq N_d(\delta, \mathcal{T}). \quad (3.34)$$

The second inequality comes from (3.33). For the sake of contradiction, if the first inequality were not true, then two elements of $\mathcal{T}_{2\delta}$ should belong to a ball centred at a point of \mathcal{T}'_δ with radius δ , which is absurd.

Example 3.49. For all $\delta > 0$, the δ -covering number of the Euclidean ball $\mathcal{B}(0, 1)$ in \mathbb{R}^d satisfies:

$$N'_{\|\cdot\|}(\delta, \mathcal{B}(0, 1)) \leq \left(\frac{3}{\delta}\right)^d.$$

Proof. We will derive an upper bound for the packing number $N_{\|\cdot\|}(\delta, \mathcal{T})$ and conclude with (3.34). If $\mathcal{T}_\delta = \{t_1, t_2, \dots, t_k\}$ is a δ -packing of $\mathcal{B}(0, 1)$ for $\delta \leq 1$, then the balls $\mathcal{B}(t_i, \frac{\delta}{2})$ are not intersecting. Thus, the Lebesgue volume of $\bigcup_{i \in \llbracket 1, k \rrbracket} \mathcal{B}(t_i, \frac{\delta}{2})$ equals $k \left(\frac{\delta}{2}\right)^d \omega_d$, with ω_d the Lebesgue volume of $\mathcal{B}(0, 1)$. Since $\bigcup_{i \in \llbracket 1, k \rrbracket} \mathcal{B}(t_i, \frac{\delta}{2}) \subset \mathcal{B}(0, \frac{3}{2})$, it comes that $k \left(\frac{\delta}{2}\right)^d \omega_d \leq \left(\frac{3}{2}\right)^d \omega_d$, which concludes. \square

We define the δ -entropy number as the logarithm of the δ -packing number:

$$H_d(\delta, \mathcal{T}) = \log(N_d(\delta, \mathcal{T})). \quad (3.35)$$

The function $H_d(\cdot, \mathcal{T})$ is called the *metric entropy* of \mathcal{T} for the pseudo-metric d .

The Dudley entropy integral

The Dudley entropy integral furnishes an upper-bound for the expectation of the maximal deviation on a finite set, depending on the metric entropy.

The theorem

Theorem 3.50 (Dudley entropy integral [BLM13, Corollary 13.2]). *Let \mathcal{T} be a finite space equipped with a semi-metric d . Let $(X_t)_{t \in \mathcal{T}}$ be a collection of random variables such that*

$$\log(\mathbb{E}[\exp(\lambda(X_t - X_{t'}))]) \leq \frac{\lambda^2 d^2(t, t')}{2}$$

for all $t, t' \in \mathcal{T}$ and all $\lambda > 0$. Then, for all $t_0 \in \mathcal{T}$:

$$\mathbb{E} \left[\sup_{t \in \mathcal{T}} X_t - X_{t_0} \right] \leq 12 \int_0^{\frac{\delta}{2}} \sqrt{H_d(u, \mathcal{T})} du,$$

where $\delta = \sup_{t \in \mathcal{T}} d(t, t_0)$.

Set \mathcal{F} a finite set of measurable functions on \mathcal{X} taking values in $[-R, R]$ for some $R > 0$. As noted in [BLM13], we can apply Theorem 3.50 to the subset of the Euclidean ball $\mathcal{B}(0, 1)$ of \mathbb{R}^n , $\mathcal{T} = \frac{1}{R\sqrt{n}}\mathcal{F}(x_1^n)$ which elements are of the form $t = (\frac{1}{R\sqrt{n}}f_t(x_i))_{i \in [1, n]}$ and the associated random variables $X_t = \frac{1}{R\sqrt{n}}\sum_{i=1}^n f_t(x_i)\sigma_i$ for the σ_i s independent Rademacher random variables, for the x_i s in \mathcal{X} . Then, the condition in Theorem 3.50 is satisfied for the Euclidean distance, that is $d(t, t') = \frac{1}{R\sqrt{n}}\|f_t(x_1^n) - f_{t'}(x_1^n)\| = \frac{1}{R}\sqrt{\frac{1}{n}\sum_{i=1}^n (f_t(x_i) - f_{t'}(x_i))^2}$, where we used the fact that $\mathbb{E}[e^{sY}] \leq \exp\left(\frac{s^2(d-c)^2}{8}\right)$ whenever Y is a random variable taking values in $[c, d]$. In particular, provided that $f_0 : x \mapsto 0$ belongs to \mathcal{F} , it holds that:

$$\tilde{R}_n(\mathcal{F}(x_1^n)) = \frac{R}{\sqrt{n}}\mathbb{E}_\sigma \left[\sup_{t \in \mathcal{T}} \frac{1}{R\sqrt{n}} \sum_{i=1}^n f_t(x_i)\sigma_i \right] \leq 12 \frac{R}{\sqrt{n}} \int_0^{\frac{1}{2}} \sqrt{H_{\|\cdot\|} \left(u, \frac{1}{R\sqrt{n}}\mathcal{F}(x_1^n) \right)} du,$$

where $\tilde{R}_n(\mathcal{F}(x_1^n))$ is defined by Equation (3.28) and $\|\cdot\|$ denotes the Euclidean norm in \mathbb{R}^n .

From (3.29), it comes that:

$$\mathbb{E} \left[\sup_{f \in \mathcal{F}} |Pf - P_n f| \right] \leq 24 \frac{R}{\sqrt{n}} \mathbb{E} \left[\int_0^{\frac{1}{2}} \sqrt{H_{\|\cdot\|} \left(u, \frac{1}{R\sqrt{n}}(\mathcal{F} \cup -\mathcal{F})(X_1^n) \right)} du \right], \quad (3.36)$$

with $-\mathcal{F} = \{-f \mid f \in \mathcal{F}\}$.

According to (3.34) and (3.35), for all $\delta > 0$, we can bound the term $H_{\|\cdot\|} \left(\delta, \frac{1}{R\sqrt{n}}(\mathcal{F} \cup -\mathcal{F})(X_1^n) \right)$ from above by $\log \left(N'_{\|\cdot\|} \left(\frac{\delta}{2}, \frac{1}{R\sqrt{n}}(\mathcal{F} \cup -\mathcal{F})(X_1^n) \right) \right)$.

Some properties about the metric entropy are derived in the following section and may be used to derive upper-bounds for the Dudley integral.

Some properties for the covering number of sets of type $\mathcal{F}(x_1^n)$ For all classes of measurable functions \mathcal{F}_1 and \mathcal{F}_2 , set $\mathcal{F}_1 + \mathcal{F}_2 = \{f_1 + f_2 \mid f_1 \in \mathcal{F}_1, f_2 \in \mathcal{F}_2\}$ and $\mathcal{F}_1 \times \mathcal{F}_2 = \{f_1.f_2 \mid f_1 \in \mathcal{F}_1, f_2 \in \mathcal{F}_2\}$. Then,

$$N'_{\|\cdot\|}(2\delta, \mathcal{F}_1 + \mathcal{F}_2(x_1^n)) \leq N'_{\|\cdot\|}(\delta, \mathcal{F}_1(x_1^n))N'_{\|\cdot\|}(\delta, \mathcal{F}_2(x_1^n)), \quad (3.37)$$

and under the assumption that the elements of \mathcal{F}_1 and \mathcal{F}_2 are all bounded by 1:

$$N'_{\|\cdot\|}(2\delta, \mathcal{F}_1 \times \mathcal{F}_2(x_1^n)) \leq N'_{\|\cdot\|}(\delta, \mathcal{F}_1(x_1^n))N'_{\|\cdot\|}(\delta, \mathcal{F}_2(x_1^n)). \quad (3.38)$$

For all class of measurable functions \mathcal{F} , if $-\mathcal{F} = \{-f \mid f \in \mathcal{F}\}$ then

$$N'_{\|\cdot\|}(\delta, \mathcal{F} \cup -\mathcal{F}(x_1^n)) \leq 2N'_{\|\cdot\|}(\delta, \mathcal{F}(x_1^n)). \quad (3.39)$$

Moreover, for all class \mathcal{F} of measurable functions, if $\mathcal{F}_{(k)} = \{\min_{l \in [1, k]} f_l \mid f_l \in \mathcal{F}\}$, then

$$N'_{\|\cdot\|}(\sqrt{k}\delta, \mathcal{F}_{(k)}(x_1^n)) \leq \left(N'_{\|\cdot\|}(\delta, \mathcal{F}(x_1^n)) \right)^k, \quad (3.40)$$

and for $\mathcal{F}_{(k)+} = \{\max_{l \in [1, k]} f_l \mid f_l \in \mathcal{F}\}$, then

$$N'_{\|\cdot\|}(\sqrt{k}\delta, \mathcal{F}_{(k)+}(x_1^n)) \leq \left(N'_{\|\cdot\|}(\delta, \mathcal{F}(x_1^n)) \right)^k. \quad (3.41)$$

Proof. For the three first inequalities, if $\bigcup_{i \in \llbracket 1, N' \rrbracket} \mathcal{B}_{\|\cdot\|}(c^i, \delta)$ covers $\mathcal{F}_1(x_1^n)$ and $\bigcup_{i \in \llbracket 1, M' \rrbracket} \mathcal{B}_{\|\cdot\|}(c'^i, \delta)$ covers $\mathcal{F}_2(x_1^n)$, then it suffices to take balls centred at elements $c^i + c'^j, c^i c'^j$ or respectively c^i and $-c^i$ to cover $\mathcal{F}_1 + \mathcal{F}_2, \mathcal{F}_1 \times \mathcal{F}_2$ and respectively $\mathcal{F}_1 \cup -\mathcal{F}_1$.

For the fourth inequality, set $c^1, c^2, \dots, c^{N'}$ in \mathbb{R}^n that satisfy $\bigcup_{j \in \llbracket 1, N' \rrbracket} \mathcal{B}_{\|\cdot\|}(c^j, \delta) \supset \mathcal{F}(x_1^n)$. Then for all $i \in \llbracket 1, n \rrbracket$, for all $f_1, f_2, \dots, f_k \in \mathcal{F}$, with $c^{(1)}, c^{(2)}, \dots, c^{(k)}$ in \mathbb{R}^n such that $f_j(x_1^n) \in \mathcal{B}_{\|\cdot\|}(c^{(j)}, \delta)$ for all $j \in \llbracket 1, k \rrbracket$, then we have:

$$\left(\min_{l \in \llbracket 1, k \rrbracket} f_l(x_i) - \min_{l \in \llbracket 1, k \rrbracket} c^{(l)}(x_i) \right)^2 \leq \sum_{l=1}^k \left(f_l(x_i) - c^{(l)}(x_i) \right)^2.$$

Also note that the number of elements of type $\min_{l \in \llbracket 1, k \rrbracket} c^{(l)}(x_1^n)$ is at most $\binom{N'}{k} \leq N'^k$. Equation (3.40) follows. The last inequality comes from Equation (3.40) after noting that $\max_{i \in \llbracket 1, k \rrbracket} f_i = -\min_{i \in \llbracket 1, k \rrbracket} (-f_i)$. \square

It is also possible to derive bounds for families \mathcal{F} indexed by a set \mathcal{T} equipped with a distance d , under Lipschitz assumption of type $|f_t(x) - f_{t'}(x)| \leq Cd(t, t')$ for all t, t' in \mathcal{T} and all $x \in (x_1^n)$. Under such an assumption, it holds that

$$N'_{\|\cdot\|} \left(\delta, \frac{1}{\sqrt{n}} \mathcal{F}(x_1^n) \right) \leq N'_d \left(\frac{\delta}{2C}, \mathcal{T} \right). \quad (3.42)$$

Such a bound comes from the following lemma.

Lemma 3.51. *If (\mathcal{T}, d) and (\mathcal{T}', d') with $\mathcal{T}' = f(\mathcal{T}) = \{f(t) \mid t \in \mathcal{T}\}$ are two metric spaces for some C -Lipschitz function $f : \mathcal{T} \mapsto \mathcal{T}'$, with $C > 0$, that is : $d(f(t_1), f(t_2)) \leq Cd(t_1, t_2)$. Then,*

$$N'_{d'}(\delta, \mathcal{T}') \leq N'_d \left(\frac{\delta}{2C}, \mathcal{T} \right).$$

Proof. For $c_1, c_2, \dots, c_{N'}$ such that $\mathcal{T} \subset \bigcup_{i \in \llbracket 1, N' \rrbracket} \mathcal{B}_d(c_i, \frac{\delta}{2C})$, set $t_i \in \mathcal{T}$ such that $d(c_i, t_i) \leq \frac{\delta}{2C}$. Then, for $t \in \mathcal{T} \cap \mathcal{B}_d(c_i, \frac{\delta}{2C})$, then $f(t) \in \mathcal{B}_{d'}(f(t_i), \delta)$, which concludes. \square

Shattering dimension to bound metric entropy

For \mathcal{F} , a class of measurable real-valued functions bounded by 1, an upper-bound relying on a combinatorial quantity, the shattering dimension, can be derived for the covering number $N'_{\|\cdot\|}(\delta, \mathcal{F}(x_1^n))$, for all $x_1^n \in \mathcal{X}^n$ and $\delta \in (0, 1)$.

Definition of shattering dimension – The link with the Vapnik-Chervonenkis dimension

Given some $\delta > 0$, we say that a subset S of \mathcal{X} is δ -shattered by \mathcal{F} if there exists some function h on S such that:

$$\forall S' \subset S, \exists f \in \mathcal{F}, \text{ s.t. } \forall x \in S', f(x) \leq h(x) \text{ and } \forall x \in S \setminus S', f(x) \geq h(x) + \delta. \quad (3.43)$$

We denote by $V(\mathcal{F}, \delta)$ the *shattering dimension* of \mathcal{F} , which is the maximal cardinality of a set S δ -shattered by \mathcal{F} .

Note that the shattering dimension of a set of indicator functions \mathcal{F} coincides with the maximal shattering dimension of $\mathcal{F}(x_1^n)$, defined in Section 3.A.2, in the sense that for all $\delta \in (0, 1)$:

$$V(\mathcal{F}, \delta) = \sup_{n \in \mathbb{N}, x_1^n \in \mathcal{X}^n} V(\mathcal{F}(x_1^n)). \quad (3.44)$$

Examples of computations of shattering dimension The two following lemmas provide a computation of the shattering dimension of the set of all indicators of balls, and the set of all functions $x \mapsto \frac{1}{R} \mathbb{1}_{\mathcal{B}(0,R)} \langle x, v \rangle$ for all $v \in S(0, 1)$ on \mathbb{R}^d .

Lemma 3.52. Set $\mathcal{F} = \{ \mathbb{1}_{\mathcal{B}(x,r)} \mid x \in \mathbb{R}^d, r > 0 \}$ for some $R > 0$, then, for all $\delta \in (0, 1)$,

$$V(\mathcal{F}, \delta) \leq d + 1.$$

Proof. Set S a subset of \mathbb{R}^d which is δ -shattered by \mathcal{F} , then necessarily, for all subset S' of S , there is some $x \in \mathbb{R}^d$ and $r > 0$ such that $y \notin \mathcal{B}(x, r)$ if $y \in S'$ and $y \in \mathcal{B}(x, r)$ if $y \in S \setminus S'$. According to [Dud79], this is possible only if $|S| \leq d + 1$. Hence, $V(\mathcal{F}, \delta) \leq d + 1$. \square

Lemma 3.53. Set $\mathcal{F} = \{ x \mapsto \frac{1}{R} \mathbb{1}_{\mathcal{B}(0,R)}(x) \langle x, v \rangle \mid v \in S(0, 1) \}$ for $R > 0$, a family of functions defined on \mathbb{R}^d , then for all $\delta \in (0, 1)$,

$$V(\mathcal{F}, \delta) \leq d + 2.$$

Proof. Set $S = \{x_1, x_2, \dots, x_V\}$ a subset of $\mathcal{B}(0, R)$ which is δ -shattered by \mathcal{F} , and the function h as defined in (3.43). Then for all $S' \subset S$, there is some $v \in S(0, 1)$ such that $\langle (x, -h(x)), (v, 1) \rangle \leq 0$ if $x \in S'$ and $\langle (x, -h(x)), (v, 1) \rangle \geq \delta$ if $x \in S \setminus S'$. In particular, it means that all subset $(S', -h(S'))$ rewrites as the intersection of $(S, -h(S))$ with a half space in \mathbb{R}^{d+2} . According to [Dud79], this is possible only if $|S| \leq d + 2$. Hence, $V(\mathcal{F}, \delta) \leq d + 2$. \square

Shattering dimension and covering number For P' , a probability measure on \mathcal{X} , for $\delta > 0$, we denote by $N(\mathcal{F}, \delta, L_2(P'))$ the minimal cardinality of a set of functions \mathcal{F}_δ on \mathcal{X} needed to cover \mathcal{F} within a $L_2(P')$ -error equal to δ . That is such that

$$\forall f \in \mathcal{F}, \exists f' \in \mathcal{F}_\delta, \text{ s.t. } \sqrt{\int_{\mathcal{X}} (f(u) - f'(u))^2 dP'(u)} \leq \delta.$$

Theorem 3.54 ([MV03, Theorem 1]). Let \mathcal{F} be a class of functions bounded by 1 defined on a set \mathcal{X} . Then, for all measure P' in \mathcal{X} and $\delta \in (0, 1)$:

$$N(\mathcal{F}, \delta, L_2(P')) \leq \left(\frac{2}{\delta} \right)^{KV(\mathcal{F}, c\delta)},$$

for K and c some positive absolute constants.

In particular, an application of Theorem 3.54 to all possible measures $P' = P_n$ yields that for all class \mathcal{F} of functions bounded by 1 on \mathcal{X} , for all $x_1^n \in \mathcal{X}^n$ and $\delta \in (0, 1)$, for some absolute constants K and c that:

$$N'_{\|\cdot\|} \left(\delta, \frac{1}{\sqrt{n}} \mathcal{F}(x_1^n) \right) \leq \left(\frac{2}{\delta} \right)^{KV(\mathcal{F}, c\delta)}, \tag{3.45}$$

where we recall that $N'_{\|\cdot\|}(\delta, \mathcal{F}(x_1^n))$ is the δ -covering number of $\mathcal{F}(x_1^n)$ for the Euclidean norm on \mathbb{R}^n , that is:

$$\|f(x_1^n) - f'(x_1^n)\| = \sqrt{\sum_{i=1}^n (f(x_i) - f'(x_i))^2}.$$

A ready-to-use theorem

As a consequence, we get an upper bound for the expectation of the maximal deviation in terms of the shattering dimension of \mathcal{F} :

Theorem 3.55. *For all countable family \mathcal{F} of measurable functions taking values in $[-R, R]$ including the function $x \mapsto 0$, it holds:*

$$\begin{aligned} \mathbb{E} \left[\sup_{f \in \mathcal{F}} |Pf - P_n f| \right] &\leq 24 \frac{R}{\sqrt{n}} \mathbb{E} \left[\int_0^{\frac{1}{2}} \sqrt{\log \left(N'_{\|\cdot\|} \left(\frac{u}{2}, \frac{1}{R\sqrt{n}} (\mathcal{F} \cup -\mathcal{F})(X_1^n) \right) \right)} du \right] \\ &\leq 24 \frac{R}{\sqrt{n}} \int_0^{\frac{1}{2}} \sqrt{KV(c) \log \left(\frac{4}{u} \right)} du \end{aligned}$$

for some absolute positive constants c and K , with $V(c) = \sup_{u \in (0, \frac{1}{2})} V \left(\frac{1}{R} (\mathcal{F} \cup -\mathcal{F}), cu \right)$ the maximal shattering dimension of $\frac{1}{R} (\mathcal{F} \cup -\mathcal{F})$.

Proof. It suffices to apply Equation 3.45 to $\frac{1}{R\sqrt{n}} (\mathcal{F} \cup -\mathcal{F})(x_1^n)$ together with Equation 3.36. To go from a finite family \mathcal{F} to a countable family, it suffices to apply the dominated convergence lemma, and to note that when $\mathcal{F} \subset \mathcal{F}'$, then for all $\delta > 0$, $V(\mathcal{F}, \delta) \leq V(\mathcal{F}', \delta)$. \square

Some applications

It appends that the class of functions $\frac{1}{R\sqrt{n}} (\mathcal{F} \cup -\mathcal{F})$ rewrites as a product, sum, minimum or maximum of more elementary classes. Then, it would be appropriate to apply the first inequality of Theorem 3.55 together with the properties in Section 3.A.3, and then to apply Theorem 3.54 to bound the covering numbers by a quantity involving the shattering dimension.

The proofs of the following Corollaries are based on this method.

Bounds for truncated moments of order 1 and order 2 moments in Euclidean \mathbb{R}^d

Corollary 3.56. *If $P \in \mathcal{P}^{(V)}(\mathbb{R}^d)$ is a sub-Gaussian measure for some $V > 0$, then for all $p > 0$, with probability larger than $1 - 3n^{-p}$, it holds*

$$\sup_{x \in \mathbb{R}^d, r > 0} \|(P_n - P)u \mathbb{1}_{\mathcal{B}(x,r)}(u)\| \leq CV\sqrt{d+2} \frac{(1+p) \log(n)}{\sqrt{n}}.$$

and

$$\sup_{x \in \mathbb{R}^d, r > 0} \|(P_n - P)u \mathbb{1}_{\overline{\mathcal{B}}(x,r)}(u)\| \leq CV\sqrt{d+2} \frac{(1+p) \log(n)}{\sqrt{n}}.$$

In particular,

$$\sup_{x \in \mathbb{R}^d, r > 0} \|(P_n - P)u \mathbb{1}_{\partial \mathcal{B}(x,r)}(u)\| \leq 2CV\sqrt{d+2} \frac{(1+p) \log(n)}{\sqrt{n}}.$$

The same holds for half-spaces.

Proof. First note that since $\mathbb{R}^d \times \mathbb{R}_+$ is separable, we can find a dense and countable subset \mathcal{D} of $\mathbb{R}^d \times \mathbb{R}_+$. According to the dominated convergence lemma and the fact that $\mathcal{B}(x, r)$ can be approximated from below with $\mathcal{B}(x_N, r_N)$ with $(x_N, r_N) \in \mathcal{D}$, the supremum on \mathcal{D} and $\mathbb{R}^d \times \mathbb{R}_+$ of $\|(P_n - P)u \mathbb{1}_{\mathcal{B}(x,r)}(u)\|$ coincide. Thus there is no problem of measurability.

Let $\lambda = p \log(n)$, $R = \sqrt{4V^2(\log(n) + \lambda)}$. We may write

$$\begin{aligned} \sup_{x,r} \|(P_n - P)u \mathbb{1}_{\mathcal{B}(x,r)}(u)\| &= \sup_{x,r} \left\| \frac{1}{n} \sum_{i=1}^n X_i \mathbb{1}_{\mathcal{B}(x,r)}(X_i) - Pu \mathbb{1}_{\mathcal{B}(x,r)}(u) \right\| \\ &\leq \sup_{x,r} \left\| \frac{1}{n} \sum_{i=1}^n X_i \mathbb{1}_{\mathcal{B}(x,r)}(X_i) \mathbb{1}_{\|X_i\| \leq R} - Pu \mathbb{1}_{\mathcal{B}(x,r)}(u) \mathbb{1}_{\|u\| \leq R} \right\| \\ &\quad + P \|u\| \mathbb{1}_{\|u\| > R} + \frac{1}{n} \sum_{i=1}^n \|X_i\| \mathbb{1}_{\|X_i\| > R}. \end{aligned}$$

From Lemma 4.36,

$$P(\|u\| \mathbb{1}_{\|u\| > R}) \leq 2Vn^{-(p+1)}$$

and with probability larger than $1 - n^{-2p-1}$, it holds

$$\frac{1}{n} \sum_{i=1}^n \|X_i\| \mathbb{1}_{\|X_i\| > R} = 0.$$

As a consequence of the Cauchy-Schwarz inequality, we may write

$$\sup_{x \in \mathbb{R}^d, r > 0} \|(P_n - P)u \mathbb{1}_{\mathcal{B}(x,r)}(u) \mathbb{1}_{\mathcal{B}(0,R)}(u)\| = \sup_{x \in \mathbb{R}^d, r > 0, v \in \mathbb{S}(0,1)} |(P_n - P)f_{x,r,v}(u)|,$$

where $f_{x,r,v}(u) = \langle v, u \mathbb{1}_{\mathcal{B}(x,r)}(u) \mathbb{1}_{\mathcal{B}(0,R)}(u) \rangle$ for all $u \in \mathbb{R}^d$.

According to (3.25) with $\delta = n^{-p}$, with probability larger than $1 - n^{-p}$ it holds

$$\sup_{x \in \mathbb{R}^d, r > 0, v \in \mathbb{S}(0,1)} |P_n - P|f_{x,r,v}(u)| \leq \mathbb{E} \left[\sup_{x \in \mathbb{R}^d, r > 0, v \in \mathbb{S}(0,1)} |P_n - P|f_{x,r,v}(u)| \right] + \sqrt{\frac{2R^2 p \log(n)}{n}}.$$

It remains to derive an upper-bound for $\mathbb{E} \left[\sup_{x \in \mathbb{R}^d, r > 0, v \in \mathbb{S}(0,1)} |P_n - P|f_{x,r,v}(u)| \right]$.

According to Theorem 3.55

$$\mathbb{E} \left[\sup_{x \in \mathbb{R}^d, r > 0, v \in \mathbb{S}(0,1)} |P_n - P|f_{x,r,v}(u)| \right] \leq 24 \frac{R}{\sqrt{n}} \mathbb{E} \left[\int_0^{\frac{1}{2}} \sqrt{\log \left(N'_{\|\cdot\|} \left(\frac{u}{2}, \frac{1}{R\sqrt{n}} \mathcal{F}(X_1^n) \right) \right)} du \right]$$

with $\frac{1}{R} \mathcal{F}(X_1^n) = \mathcal{F}_1 \times \mathcal{F}_2(X_1^n)$ for \mathcal{F}_1 the set of indicators of balls and \mathcal{F}_2 the set of functions $x \mapsto \frac{1}{R} \mathbb{1}_{\mathcal{B}(0,R)} \langle x, v \rangle$ for all $v \in \mathbb{S}(0,1)$. According to 3.38, for all $\delta > 0$,

$$\begin{aligned} N'_{\|\cdot\|} \left(\frac{u}{2}, \frac{1}{\sqrt{n}} \mathcal{F}_1 \times \mathcal{F}_2(x_1^n) \right) &= N'_{\|\cdot\|} \left(\frac{u}{2} \sqrt{n}, \mathcal{F}_1 \times \mathcal{F}_2(x_1^n) \right) \\ &\leq N'_{\|\cdot\|} \left(\frac{u\sqrt{n}}{4}, \mathcal{F}_1(x_1^n) \right) N'_{\|\cdot\|} \left(\frac{u\sqrt{n}}{4}, \mathcal{F}_2(x_1^n) \right) \\ &\leq N'_{\|\cdot\|} \left(\frac{u}{4}, \frac{1}{\sqrt{n}} \mathcal{F}_1(x_1^n) \right) N'_{\|\cdot\|} \left(\frac{u}{4}, \frac{1}{\sqrt{n}} \mathcal{F}_2(x_1^n) \right). \end{aligned}$$

From Lemma 3.52, $V(\mathcal{F}_1, \delta) \leq d + 1$, thus from (3.45),

$$N'_{\|\cdot\|} \left(\frac{u}{4}, \frac{1}{\sqrt{n}} \mathcal{F}_1(x_1^n) \right) \leq \left(\frac{8}{u} \right)^{K(d+1)}.$$

From Lemma 3.53, $V(\mathcal{F}_2, \delta) \leq d + 2$, thus

$$N'_{\|\cdot\|}\left(\frac{u}{4}, \frac{1}{\sqrt{n}}\mathcal{F}_2(x_1^n)\right) \leq \left(\frac{8}{u}\right)^{K(d+2)}.$$

As a consequence,

$$N'_{\|\cdot\|}\left(\frac{u}{2}, \frac{1}{R\sqrt{n}}\mathcal{F}(x_1^n)\right) \leq \left(\frac{8}{u}\right)^{K(2d+3)}.$$

Thus, putting all bounds together, with probability larger than $1 - 2n^{-p}$ we have:

$$\begin{aligned} \sup_{x \in \mathbb{R}^d, r > 0, v \in \mathcal{S}(0,1)} |P_n - P|f_{x,r,v}(u) &\leq CR \frac{\sqrt{2d+3}}{\sqrt{n}} + 2Vn^{-(p+1)} + \sqrt{\frac{2R^2p \log(n)}{n}} \\ &\leq CV\sqrt{d+2} \frac{(1+p) \log(n)}{\sqrt{n}} \end{aligned}$$

for some absolute positive constant C . □

It is also possible to derive an upper-bound for the deviation of moments of order 2 restricted to balls, as follows.

Corollary 3.57. *Let $P \in \mathcal{P}^{(V)}(\mathbb{R}^d)$ be a sub-Gaussian measure on \mathbb{R}^d with variance V^2 . Then, with probability larger than $1 - 3n^{-p}$,*

$$\begin{aligned} \sup_{x \in \mathbb{R}^d, r > 0} |P\|u\|^2 \mathbb{1}_{\mathcal{B}(x,r)}(u) - P_n\|u\|^2 \mathbb{1}_{\mathcal{B}(x,r)}(u)| &\leq CV^2\sqrt{d+1} \frac{((1+p) \log(n))^{\frac{3}{2}}}{\sqrt{n}}, \\ \sup_{x \in \mathbb{R}^d, r > 0} |P\|u\|^2 \mathbb{1}_{\overline{\mathcal{B}}(x,r)}(u) - P_n\|u\|^2 \mathbb{1}_{\overline{\mathcal{B}}(x,r)}(u)| &\leq CV^2\sqrt{d+1} \frac{((1+p) \log(n))^{\frac{3}{2}}}{\sqrt{n}}, \end{aligned}$$

in particular,

$$\sup_{x \in \mathbb{R}^d, r > 0} |P\|u\|^2 \mathbb{1}_{\partial\mathcal{B}(x,r)}(u) - P_n\|u\|^2 \mathbb{1}_{\partial\mathcal{B}(x,r)}(u)| \leq 2CV^2\sqrt{d+1} \frac{((1+p) \log(n))^{\frac{3}{2}}}{\sqrt{n}}.$$

The same holds for half-spaces.

Proof. For the same reason as Corollary 3.56, the supremum of $|P\|u\|^2 \mathbb{1}_{\mathcal{B}(x,r)}(u) - P_n\|u\|^2 \mathbb{1}_{\mathcal{B}(x,r)}(u)|$ on $\mathbb{R}^d \times \mathbb{R}_+$ coincides with the supremum taken on \mathcal{D} a dense and countable subset of $\mathbb{R}^d \times \mathbb{R}_+$, thus there is no problem of measurability.

Let $\lambda = p \log(n)$ and $R = \sqrt{4V^2(\log(n) + \lambda)}$. We may write

$$\begin{aligned} \sup_{x,r} |(P_n - P)\|u\|^2 \mathbb{1}_{\mathcal{B}(x,r)}(u)| &= \sup_{x,r} \left| \frac{1}{n} \sum_{i=1}^n \|X_i\|^2 \mathbb{1}_{\mathcal{B}(x,r)}(X_i) - P(\|u\|^2 \mathbb{1}_{\mathcal{B}(x,r)}(u)) \right| \\ &\leq \sup_{x,r} \left| \frac{1}{n} \sum_{i=1}^n \|X_i\|^2 \mathbb{1}_{\mathcal{B}(x,r)}(X_i) \mathbb{1}_{\|X_i\| \leq R} - P(\|u\|^2 \mathbb{1}_{\mathcal{B}(x,r)}(u) \mathbb{1}_{\|u\| \leq R}) \right| \\ &\quad + P(\|u\|^2 \mathbb{1}_{\|u\| > R}) + \frac{1}{n} \sum_{i=1}^n \|X_i\|^2 \mathbb{1}_{\|X_i\| > R}. \end{aligned}$$

According to Lemma 4.36,

$$P(\|u\|^2 \mathbb{1}_{\|u\|>R}) \leq 3V^2 n^{-(p+1)}$$

and, with probability larger than $1 - n^{-1-2p}$,

$$\frac{1}{n} \sum_{i=1}^n \|X_i\|^2 \mathbb{1}_{\|X_i\|>R} = 0.$$

Then, according to (3.24) with $\delta = n^{-p}$, with probability larger than n^{-p} it holds

$$\begin{aligned} & \sup_{x,r} \left| \frac{1}{n} \sum_{i=1}^n \|X_i\|^2 \mathbb{1}_{\mathcal{B}(x,r)}(X_i) \mathbb{1}_{\mathcal{B}(0,R)}(X_i) - P(\|u\|^2 \mathbb{1}_{\mathcal{B}(x,r)}(u) \mathbb{1}_{\mathcal{B}(0,R)}(u)) \right| \leq \\ & \mathbb{E} \left[\sup_{x,r} \left| \frac{1}{n} \sum_{i=1}^n \|X_i\|^2 \mathbb{1}_{\mathcal{B}(x,r)}(X_i) \mathbb{1}_{\mathcal{B}(0,R)}(X_i) - P(\|u\|^2 \mathbb{1}_{\mathcal{B}(x,r)}(u) \mathbb{1}_{\mathcal{B}(0,R)}(u)) \right| \right] + \sqrt{\frac{R^4 p \log(n)}{2n}}. \end{aligned}$$

According to Theorem 3.55

$$\begin{aligned} & \mathbb{E} \left[\sup_{x,r} \left| \frac{1}{n} \sum_{i=1}^n \|X_i\|^2 \mathbb{1}_{\mathcal{B}(x,r)}(X_i) \mathbb{1}_{\mathcal{B}(0,R)}(X_i) - P(\|u\|^2 \mathbb{1}_{\mathcal{B}(x,r)}(u) \mathbb{1}_{\mathcal{B}(0,R)}(u)) \right| \right] \\ & \leq 24 \frac{R^2}{\sqrt{n}} \mathbb{E} \left[\int_0^{\frac{1}{2}} \sqrt{\log \left(N'_{\|\cdot\|} \left(\frac{u}{2}, \frac{1}{R^2 \sqrt{n}} (\mathcal{F} \cup -\mathcal{F})(X_1^n) \right) \right)} du \right] \end{aligned}$$

with $\frac{1}{R^2} (\mathcal{F} \cup -\mathcal{F})(X_1^n) = (\mathcal{F}_1 \cup -\mathcal{F}_1) \times \mathcal{F}_2(X_1^n)$ with \mathcal{F}_1 is the set of indicators of balls and \mathcal{F}_2 the single function $\{x \mapsto \frac{\|x\|^2}{R^2} \mathbb{1}_{\mathcal{B}(0,R)}(x)\}$. According to 3.38, for all $\delta > 0$,

$$N'_{\|\cdot\|} \left(2\delta, \frac{1}{R^2 \sqrt{n}} (\mathcal{F}_1 \cup -\mathcal{F}_1) \times \mathcal{F}_2(x_1^n) \right) \leq N'_{\|\cdot\|} \left(\delta, \frac{1}{\sqrt{n}} (\mathcal{F}_1 \cup -\mathcal{F}_1)(x_1^n) \right) N'_{\|\cdot\|} \left(\delta, \frac{1}{\sqrt{n}} \mathcal{F}_2(x_1^n) \right)$$

and according to (3.39),

$$N'_{\|\cdot\|} \left(\delta, \frac{1}{\sqrt{n}} (\mathcal{F}_1 \cup -\mathcal{F}_1)(x_1^n) \right) \leq 2N'_{\|\cdot\|} \left(\delta, \frac{1}{\sqrt{n}} \mathcal{F}_1(x_1^n) \right).$$

Moreover, from (3.45)

$$N'_{\|\cdot\|} \left(\delta, \frac{1}{\sqrt{n}} \mathcal{F}_1(x_1^n) \right) \leq \left(\frac{2}{\delta} \right)^{KV(\mathcal{F}_1, c\delta)}$$

and since the VC dimension of indicators of balls is $d + 1$ from [Dud79], see Lemma 3.52, it comes:

$$N'_{\|\cdot\|} \left(\delta, \frac{1}{\sqrt{n}} (\mathcal{F} \cup -\mathcal{F})(x_1^n) \right) \leq 2 \left(\frac{2}{\delta} \right)^{K(d+1)}$$

for some positive constant K .

Thus, after putting all bounds together it comes that with probability larger than $1 - 2n^{-p}$,

$$\begin{aligned} & \sup_{x \in \mathbb{R}^d, r > 0} |P\|u\|^2 \mathbb{1}_{\mathcal{B}(x,r)}(u) - P_n\|u\|^2 \mathbb{1}_{\mathcal{B}(x,r)}(u)| \\ & \leq \sqrt{d+1} \frac{R^2}{\sqrt{n}} C + R^2 \sqrt{\frac{p \log(n)}{2n}} + 3V^2 n^{-p-1} \\ & \leq CV^2 \sqrt{d+1} \frac{((1+p) \log(n))^{\frac{3}{2}}}{\sqrt{n}} \end{aligned}$$

for some absolute positive constant C . The same holds for closed balls and half-spaces. \square

As a direct consequence, from the proofs of Corollary 3.56, 3.57 and 3.48 the following lemma is proved.

Lemma 3.58. *Suppose that $P \in \mathcal{P}^{(V)}(\mathbb{R}^d)$ for some $V > 0$. Then, for every $p > 0$, with probability larger than $1 - 7n^{-p}$, we have,*

$$\begin{aligned} \sup_{x,r} |(P_n - P)\mathbb{1}_{\mathcal{B}(x,r)}(u)| &\leq C\sqrt{d+1} \frac{\sqrt{p \log(n)}}{\sqrt{n}} \\ \sup_{v,c} |(P_n - P)\mathbb{1}_{\langle v,u \rangle > c}| &\leq C\sqrt{d+1} \frac{\sqrt{p \log(n)}}{\sqrt{n}} \\ \sup_{x,r} \|(P_n - P)u\mathbb{1}_{\mathcal{B}(x,r)}(u)\| &\leq CV\sqrt{d+2} \frac{(p+1) \log(n)}{\sqrt{n}} \\ \sup_{v,c} \|(P_n - P)u\mathbb{1}_{\langle v,u \rangle > c}\| &\leq CV\sqrt{d+2} \frac{(p+1) \log(n)}{\sqrt{n}} \\ \sup_{x,r} |(P_n - P)\|u\|^2\mathbb{1}_{\mathcal{B}(x,r)}(u)| &\leq CV^2\sqrt{d+1} \frac{((p+1) \log(n))^{\frac{3}{2}}}{\sqrt{n}} \\ \sup_{v,c} |(P_n - P)\|u\|^2\mathbb{1}_{\langle v,u \rangle > c}| &\leq CV^2\sqrt{d+1} \frac{((p+1) \log(n))^{\frac{3}{2}}}{\sqrt{n}}, \end{aligned}$$

where $C > 0$ denotes a universal constant, $x \in \mathbb{R}^d$, $v \in S(0, 1)$, $r > 0$ and $c \in \mathbb{R}$.

Bounds for the Bregman balls indicators and truncated Bregman-moments We can also derive deviation bounds for indicators of Bregman-balls and for the Bregman divergence, as describes by the two following lemmas.

Lemma 3.59. *Let ϕ be a differentiable convex function on $\Omega \subset \mathbb{R}^d$ and P a Borel probability distribution on Ω . Then, for every $p > 0$, with probability larger than $\exp(-p)$, it holds for some absolute positive constant C , that:*

$$\sup_{c \in \Omega^{(k)}, r \geq 0} |(P_n - P)\mathbb{1}_{\mathcal{B}_\phi(c,r)}(u)| \leq C \frac{\sqrt{k}\sqrt{d+1}}{\sqrt{n}} + \sqrt{\frac{p}{2n}}.$$

Proof. According to (3.24) with $\delta = \exp(-p)$ for some $p > 0$, with probability larger than $\exp(-p)$ it holds

$$\sup_{c \in \Omega^{(k)}, r \geq 0} |(P_n - P)\mathbb{1}_{\mathcal{B}_\phi(c,r)}(u)| \leq \mathbb{E} \left[\sup_{c \in \Omega^{(k)}, r \geq 0} |(P_n - P)\mathbb{1}_{\mathcal{B}_\phi(c,r)}(u)| \right] + \sqrt{\frac{p}{2n}}.$$

According to Theorem 3.55

$$\begin{aligned} &\mathbb{E} \left[\sup_{c \in \Omega^{(k)}, r \geq 0} |(P_n - P)\mathbb{1}_{\mathcal{B}_\phi(c,r)}(u)| \right] \\ &\leq 24 \frac{1}{\sqrt{n}} \mathbb{E} \left[\int_0^{\frac{1}{2}} \sqrt{\log \left(N'_{\|\cdot\|} \left(\frac{u}{2}, \frac{1}{\sqrt{n}} (\mathcal{F}_{(k)+} \cup -\mathcal{F}_{(k)+})(X_1^n) \right) \right)} du \right] \end{aligned}$$

with $\mathcal{F}_{(k)+} = \{\max_{l \in [1,k]} f_l \mid f_l \in \mathcal{F}\}$ with \mathcal{F} the set of indicators of Bregman-balls $\mathcal{B}_\phi(c, r)$ for $c \in \Omega$ and $r \geq 0$, thus from (3.41),

$$N'_{\|\cdot\|} \left(\frac{u}{2}, \frac{1}{\sqrt{n}} \mathcal{F}_{(k)+}(x_1^n) \right) \leq \left(N'_{\|\cdot\|} \left(\frac{u}{2\sqrt{k}}, \frac{1}{\sqrt{n}} \mathcal{F}(x_1^n) \right) \right)^k.$$

Moreover, from (3.45)

$$N'_{\|\cdot\|} \left(\delta, \frac{1}{\sqrt{n}} \mathcal{F}(x_1^n) \right) \leq \left(\frac{2}{\delta} \right)^{KV(\mathcal{F}, c\delta)}$$

for some positive absolute constants K and c . Since the VC dimension of indicators of Bregman-balls is $d + 1$, according to Proposition 3.47, with the same argumentation as in the proof of Lemma 3.52, it holds:

$$N'_{\|\cdot\|} \left(\frac{\delta}{2}, \frac{1}{\sqrt{n}} \mathcal{F}_{(k)+} \cup -\mathcal{F}_{(k)+}(x_1^n) \right) \leq 2 \left(\frac{4\sqrt{k}}{\delta} \right)^{K(d+1)k}.$$

We also used (3.39).

Thus, for every $p > 0$, with probability larger than $\exp(-p)$, it holds for some absolute positive constant C , that:

$$\sup_{\mathbf{c} \in \Omega^{(k)}, r \geq 0} \left| (P_n - P) \mathbb{1}_{\mathcal{B}_\phi(\mathbf{c}, r)}(u) \right| \leq C \frac{\sqrt{k}\sqrt{d+1}}{\sqrt{n}} + \sqrt{\frac{p}{2n}}.$$

□

Lemma 3.60. *Let ϕ be a C^2 convex function on $\Omega \subset \mathbb{R}^d$ and P a Borel probability distribution supported on a compact subset of Ω° , with finite first moment $P\|u\| = M_1$. Set $K > 0$, then for every $p > 0$, with probability larger than $\exp(-p)$, it holds for some positive constant C_{K, r^+, M_2} , that:*

$$\sup_{\mathbf{c} \in (\mathcal{B}(0, K) \cap F_0)^{(k)}, r \in [0, r^+]} \left| (P_n - P) d_\phi(u, \mathbf{c}) \mathbb{1}_{\mathcal{B}_\phi(\mathbf{c}, r)}(u) \right| \leq C_{K, r^+, M_2} (r^+)^2 \frac{\sqrt{k}\sqrt{d+1}}{\sqrt{n}} + (r^+)^2 \sqrt{\frac{p}{2n}},$$

where $F_0 = \text{Conv}(\text{Supp}(P))$.

Proof. According to (3.24) with $\delta = \exp(-p)$ for some $p > 0$, with probability larger than $\exp(-p)$ it holds

$$\begin{aligned} & \sup_{\mathbf{c} \in (\mathcal{B}(0, K) \cap F_0)^{(k)}, r \in [0, r^+]} \left| (P_n - P) d_\phi(u, \mathbf{c}) \mathbb{1}_{\mathcal{B}_\phi(\mathbf{c}, r)}(u) \right| \\ & \leq \mathbb{E} \left[\sup_{\mathbf{c} \in (\mathcal{B}(0, K) \cap F_0)^{(k)}, r \in [0, (r^+)^2]} \left| (P_n - P) d_\phi(u, \mathbf{c}) \mathbb{1}_{\mathcal{B}_\phi(\mathbf{c}, r)}(u) \right| \right] + (r^+)^2 \sqrt{\frac{p}{2n}}. \end{aligned}$$

According to Theorem 3.55, since for $x \in \mathcal{B}_\phi(\mathbf{c}, r)$, $d_\phi(x, \mathbf{c}) \leq r^2$, and Jensen inequality, we may write,

$$\begin{aligned} & \mathbb{E} \left[\sup_{\mathbf{c} \in (\mathcal{B}(0, K) \cap F_0)^{(k)}, r \in [0, r^+]} \left| (P_n - P) d_\phi(u, \mathbf{c}) \mathbb{1}_{\mathcal{B}_\phi(\mathbf{c}, r)}(u) \right| \right] \\ & \leq 24 \frac{(r^+)^2}{\sqrt{n}} \mathbb{E} \left[\int_0^{\frac{1}{2}} \sqrt{\log \left(N'_{\|\cdot\|} \left(\frac{u}{2}, \frac{1}{(r^+)^2 \sqrt{n}} (\mathcal{F}' \cup -\mathcal{F}')(X_1^n) \right) \right)} du \right] \\ & \leq 24 \frac{(r^+)^2}{\sqrt{n}} \int_0^{\frac{1}{2}} \sqrt{\mathbb{E} \left[\log \left(N'_{\|\cdot\|} \left(\frac{u}{2}, \frac{1}{(r^+)^2 \sqrt{n}} (\mathcal{F}' \cup -\mathcal{F}')(X_1^n) \right) \right) \right]} du \end{aligned}$$

with $\frac{1}{(r^+)^2}\mathcal{F}' = \mathcal{G} \times \mathcal{F}_{(k)^+}$ with $\mathcal{G} = \left\{ \frac{d_\phi(\cdot, \mathbf{c})}{(r^+)^2} \right\}$ and $\mathcal{F}_{(k)^+} = \{ \max_{l \in \llbracket 1, k \rrbracket} f_l \mid f_l \in \mathcal{F} \}$ with \mathcal{F} the set of indicators of Bregman-balls $\mathcal{B}_\phi(x, r)$ for $x \in \mathbb{R}^d$ and $r \geq 0$. From (3.38) and (3.39), for all $\delta > 0$,

$$N'_{\|\cdot\|} \left(\delta, \frac{1}{(r^+)^2\sqrt{n}} (\mathcal{F}' \cup -\mathcal{F}') (X_1^n) \right) \leq 2N'_{\|\cdot\|} \left(\frac{\delta}{2}, \frac{1}{\sqrt{n}} \mathcal{G} (X_1^n) \right) \times N'_{\|\cdot\|} \left(\frac{\delta}{2}, \frac{1}{\sqrt{n}} \mathcal{F}_{(k)^+} (X_1^n) \right).$$

With the same computation as in the proof of Lemma 3.59, we have

$$N'_{\|\cdot\|} \left(\frac{\delta}{2}, \frac{1}{\sqrt{n}} \mathcal{F}_{(k)^+} (X_1^n) \right) \leq \left(\frac{4\sqrt{k}}{\delta} \right)^{K(d+1)k},$$

for some absolute constant $K > 0$.

Now focus on \mathcal{G} . From Lemma 3.36, for some positive constant K , we have for all $x \in \Omega$ and $\mathbf{c}, \mathbf{c}' \in (F_0 \cap \mathcal{B}(0, K))^{(k)}$:

$$|d_\phi(x, \mathbf{c}) - d_\phi(x, \mathbf{c}')| \leq C_K \text{dist}(\mathbf{c}, \mathbf{c}') (1 + \|x\|).$$

Thus, from (3.42), $N'_{\|\cdot\|} \left(\delta, \frac{1}{\sqrt{n}} \mathcal{G} (X_1^n) \right) \leq N'_{\text{dist}} \left(\frac{\delta(r^+)^2}{2C_K(1 + \max_{i \in \llbracket 1, k \rrbracket} \|X_i\|)}, \mathcal{T}_k \right)$, with $N'_{\text{dist}}(\delta, \mathcal{T}_k)$, the δ -covering number of $\mathcal{T}_k = \mathcal{B}(0, K)^{(k)}$ for the metric dist defined by

$$\text{dist}(\mathbf{c}, \mathbf{c}') = \min_{\sigma \in \Sigma_k} \max_{i \in \llbracket 1, k \rrbracket} \|c_j - c'_{\sigma(j)}\|,$$

where Σ_k denotes the set of permutations of $\llbracket 1, k \rrbracket$.

Then, $N'_{\text{dist}}(\delta, \mathcal{T}_k) \leq N'_{\|\cdot\|}(\delta, \mathcal{T})^k$, where $\mathcal{T} = \mathcal{B}(0, K)$.

From Lemma 3.49, it comes that $N'_{\|\cdot\|}(\delta, \mathcal{T}) \leq \left(\frac{3K}{\delta}\right)^d$.

After putting all bounds together, it comes that

$$N'_{\|\cdot\|} \left(\delta, \frac{1}{(r^+)^2\sqrt{n}} (\mathcal{F}' \cup -\mathcal{F}') (X_1^n) \right) \leq 2 \left(\frac{12KC_K(1 + \max_{i \in \llbracket 1, k \rrbracket} \|X_i\|)}{(r^+)^2\delta} \right)^{kd} \left(\frac{4\sqrt{k}}{\delta} \right)^{K(d+1)k}.$$

Since from Jensen inequality $\mathbb{E}[\log(1 + \max_{i \in \llbracket 1, k \rrbracket} \|X_i\|)] \leq \log(kM_1)$, then for every $p > 0$, with probability larger than $\exp(-p)$, it holds for some absolute positive constant C , that:

$$\sup_{\mathbf{c} \in (\mathcal{B}(0, K) \cap F_0)^{(k)}, r \in [0, r^+]} \left| (P_n - P) d_\phi(u, \mathbf{c}) \mathbb{1}_{\mathcal{B}_\phi(\mathbf{c}, r)}(u) \right| \leq C_{K, r^+, M_1} (r^+)^2 \frac{\sqrt{k}\sqrt{d+1}}{\sqrt{n}} + (r^+)^2 \sqrt{\frac{p}{2n}}.$$

□

A new quantization method based on the distance to measure

Dans le chapitre précédent, il était question de scinder une distribution ou un nuage de points en k composantes, chacune représentée par un centre. Dans le cas d'un mélange de distributions, le paramètre intrinsèque k correspondait au nombre de composantes du mélange. Dans ce chapitre, nous traitons le cadre plus général des mesures de probabilité P dont le support \mathcal{K} (par exemple une sous-variété de \mathbb{R}^d) est compact. Notre objectif est alors d'inférer \mathcal{K} . Nous prenons le point de vue de la quantification (i.e. k grand) et développons/étudions une méthode permettant d'approcher au mieux la fonction distance au compact \mathcal{K} par une fonction distance à un ensemble de k points (centres) bien choisis. Notre méthode repose sur une approximation de la fonction distance à la mesure $d_{P,h,2}$ par une k -fonction puissance. De façon équivalente, on peut la voir comme une méthode de quantification Bregman avec une divergence de Bregman construite à partir de la distance à la mesure $d_{P,h,2}$. Aussi, comme pour le chapitre précédent, notre méthode vise à éliminer les données aberrantes lorsque l'on a accès à un nuage de points bruité autour de \mathcal{K} . Ces travaux ont été réalisés *en collaboration avec Clément Levrard*. Une version initiale de ces travaux est disponible dans le papier [BL17].

Contents

4.1	The deterministic framework	145
4.1.1	The k -PDTM: a coresets for the DTM	145
4.1.2	Proximity to the DTM	148
4.1.3	Wasserstein stability for the k -PDTM	149
4.1.4	Geometric inference with the k -PDTM	149
4.2	The statistical framework	150
4.2.1	An algorithm for the empirical k -PDTM	151
4.2.2	Proximity between the k -PDTM and its empirical version	152
4.2.3	Numerical illustrations	153
4.3	The set of local means – A strong connection with the distance-to-measure	159
4.3.1	Local means – From balls to half-spaces	160
4.3.2	Convexity of the set of local means	162
4.3.3	Duality, set of local means and distance-to-measure	163
4.4	Proofs	167
4.4.1	Proofs for Section 4.1	167
4.4.2	Proofs for Section 4.2	172
4.4.3	Proofs for Section 4.3	181

In this chapter, we consider a probability distribution P supported on a compact set \mathcal{K} (a submanifold of \mathbb{R}^d for instance). We aim to infer the function distance to the compact set \mathcal{K} . We develop/study a method to approximate the function distance to \mathcal{K} as well as possible, with a function distance to a set of k points (k large) to be determined from P . Our method is based on an approximation of the distance to the measure $d_{P,h,2}$ with a k -power distance. Equivalently, this is a quantization method based on a Bregman divergence dependent on the function distance-to-measure $d_{P,h,2}$. Our method aims at wiping out outliers or clutter noise among a set of points generated around \mathcal{K} . In this regard, the method proposed in this chapter is closely related with the work of the previous chapter.

Analyzing the sub-level sets of the distance to a compact sub-manifold of \mathbb{R}^d is a common method in TDA to understand its topology. The distance-to-measure (DTM) was introduced by Chazal, Cohen-Steiner and Mériçot in [CCSM11] to deal with the fact that the distance to a compact set is not robust against noise and outliers. This function makes the topology of

a compact subset of \mathbb{R}^d possible to infer from a noisy cloud of n points lying nearby in the Wasserstein sense. In practice, these sub-level sets may be computed using approximations of the DTM such as the q -witnessed distance [GMM11] or another power distance [BCOS15]. These approaches lead eventually to computing the homology of unions of n growing balls, which might become intractable whenever n is large.

To simultaneously deal with the two problems of large number of points and noise, we introduce the k -power distance-to-measure (k -PDTM). This new approximation of the distance-to-measure is based on a codebook of size k , optimal for some Bregman loss in the context of quantization (cf. [BMDG05]), with a data-based Bregman divergence. Each of its sub-level sets is the union of k -balls, $k \ll n$, and this distance is also proved robust to noise. We assess the quality of this approximation for k possibly dramatically smaller than n , for instance $k = n^{\frac{1}{3}}$ is proved to be optimal for 2-dimensional shapes in the next chapter.

Méridot et al noted in [GMM11] that the sublevel sets of empirical DTM are the union of around $\binom{n}{q}$ balls with $q = hn$, which makes their computation intractable in practice. To bypass this issue, approximations of the empirical DTM have been proposed in [GMM11] (q -witnessed distance) and [BCOS15] (power distance). To our knowledge, these are the only available approximations of the empirical DTM. The sublevel sets of these two approximations are the union of n balls. Thus, it makes the computation of topological invariants more tractable for small data sets, from alpha-shape for instance; see [Ede92]. Nonetheless, when n is large, there is still a need for a coreset¹ to allow one to efficiently compute an approximation of the DTM, as pointed out in [PWZ15]. In [Mér13], Méridot proves that such a coreset cannot be too small for large dimension.

Some notations and main tools

In this chapter, we work in Euclidean \mathbb{R}^d . The ball centered at c with radius r is defined by $\mathcal{B}(c, r) = \{x \in \mathbb{R}^d \mid \|x - c\| < r\}$ and its closure is denoted by $\overline{\mathcal{B}}(c, r)$. The Euclidean sphere is defined by $S(c, r) = \{x \in \mathbb{R}^d \mid \|x - c\| = r\}$. Half-spaces of direction $v \in S(0, 1)$ are denoted by $H(v, c) = \{x \in \mathbb{R}^d \mid \langle x, v \rangle > c\}$ for $c \in \mathbb{R}$, and their closure by $\overline{H}(v, c)$. For any set C , the set of all codebooks $\mathbf{c} = (c_1, c_2, \dots, c_k)$ with elements in C is denoted by $C^{(k)}$.

Also, $\mathcal{P}(\mathbb{R}^d)$ denotes the set of distributions P with support $\text{Supp}(P) \subset \mathbb{R}^d$. The subset of all distributions P in $\mathcal{P}(\mathbb{R}^d)$ with finite moment of order 2 ($P\|u\|^2 < \infty$) is denoted by $\mathcal{P}_2(\mathbb{R}^d)$. The distribution whose support is to be inferred is an element of $\mathcal{P}^K(\mathbb{R}^d) = \{P \in \mathcal{P}(\mathbb{R}^d) \mid \text{Supp}(P) \subset \overline{\mathcal{B}}(0, K)\}$ for $K > 0$. To infer $\text{Supp}(P)$, we use a modified version Q of P . This measure Q is assumed to be *sub-Gaussian* with variance $V^2 > 0$. That is, Q is a distribution in $\mathcal{P}(\mathbb{R}^d)$ such that

$$Q(\mathcal{B}(0, t)^c) \leq \exp\left(-\frac{t^2}{2V^2}\right)$$

for all $t > V$. The set of such measures is denoted by $\mathcal{P}^{(V)}(\mathbb{R}^d)$. Given $\mathbb{X}_n = \{X_1, X_2, \dots, X_n\}$ an n -sample from P , we denote by $P_n = \frac{1}{n} \sum_{i=1}^n \delta_{X_i}$ the corresponding empirical distribution.

For $P \in \mathcal{P}_2(\mathbb{R}^d)$ and $h \in (0, 1)$, we use the notation $\mathcal{P}_h(P)$ for the set of distributions $\frac{1}{h}Q \in \mathcal{P}(\mathbb{R}^d)$, with Q a sub-measure of P . The set of all of their expectations is defined by $\tilde{\mathcal{M}}_h(P) = \{m(Q) \mid Q \in \mathcal{P}_h(P)\}$, with the notation $m(Q) = Qu$ for the mean of Q , $v(Q) = Q\|u - m(Q)\|^2$ for its variance and $M(Q) = \|m(Q)\|^2 + v(Q)$ its moment of order 2. This set has nice properties given in the following lemma, whose proof is to be found in Section 4.4.1.

Lemma 4.1. *Let $P \in \mathcal{P}_2(\mathbb{R}^d)$. Then, $\tilde{\mathcal{M}}_h(P)$ is a compact convex subset of \mathbb{R}^d .*

¹A coreset for the DTM is a codebook from which we can derive a good enough approximation of the DTM.

The distributions in $\mathcal{P}_h(P)$ of interest in this chapter are given by *local distributions*. Morally, they are defined as restrictions of P on balls of P -mass h , as follows.

Definition 4.2. Let $P \in \mathcal{P}(\mathbb{R}^d)$. The *set of local distributions* at a point x , with mass parameter $h \in (0, 1)$, denoted by $\mathcal{P}_{x,h}(P)$ is the set of all distributions $P_{x,h}$ defined by $P_{x,h} = \frac{1}{h}Q$, where Q satisfies:

- (i) Q is a sub-measure of P
- (ii) $Q(\mathbb{R}^d) = h$
- (iii) Q coincides with P on $\mathcal{B}(x, \delta_{P,h}(x))$
- (iv) $\text{Supp}(Q) \subset \overline{\mathcal{B}}(x, \delta_{P,h}(x))$

Note that when P puts no mass on $\partial\mathcal{B}(x, \delta_{P,h}(x))$, the set of distributions $\mathcal{P}_{x,h}(P)$ is reduced to a singleton $\{P_{x,h}\}$ with $P_{x,h}$ defined for all Borel sets B by $P_{x,h}(B) = \frac{1}{h}P(B \cap \mathcal{B}(x, \delta_{P,h}(x)))$.

For all $P \in \mathcal{P}(\mathbb{R}^d)$, a *local mean* is defined as the expectation of a local distribution, $m(P_{x,h}) = P_{x,h}u$ associated to some $x \in \mathbb{R}^d$ and $h \in (0, 1)$. The *set of local means* of P with parameter $h \in (0, 1)$ is defined by $\mathcal{M}_h(P) = \{m(P_{x,h}) \mid x \in \mathbb{R}^d, P_{x,h} \in \mathcal{P}_{x,h}(P)\}$.

The set of local means $\mathcal{M}_h(P)$ is a subset of $\overline{\mathcal{M}}_h(P)$. A generalization to sub-measures supported on balls or half-spaces of P -mass h is given by $\overline{\mathcal{M}}_h(P)$, defined by Equation (4.10).

Brief presentation

In this chapter, we consider the quantization method of Banerjee et al. [BMDG05], with a distribution-dependent Bregman divergence. We make use of the Bregman divergence associated with the convex function² $\psi_{P,h} : x \mapsto \|x\|^2 - d_{P,h}^2(x)$. Then, the Bregman loss R is expressed in terms of local means and variances at $\mathbf{c} = (c_1, c_2, \dots, c_k)$ as

$$R(\mathbf{c}) = P \min_{i \in [1,k]} \|u - m(P_{c_i,h})\|^2 + v(P_{c_i,h}).$$

The approximation of the DTM involved in this chapter (k -PDTM) is

$$x \mapsto \min_{i \in [1,k]} \|x - m(P_{c_i^*,h})\|^2 + v(P_{c_i^*,h}),$$

for any minimizer \mathbf{c}^* of the cost function R .

This method boils down to a k -means method for which the centers are to be found among the set of local means. The penalization with the local variances forces the optimal centers to lie in areas where P has mass. Given a sample around $\mathcal{K} = \text{Supp}(P)$, this property helps to deal with outliers and clutter noise.

In section 4.1, we introduce the k -PDTM function. Different definitions are given, all proved equivalent. We derive different stability properties for this approximation of the DTM and prove that the k -PDTM is well-suited to infer $\text{Supp}(P)$, from a deterministic point of view. In section 4.2, we consider the problem within a statistical framework. We derive rates of convergence for the approximation of the DTM with the k -PDTM. This method is implemented via an algorithm inspired by the algorithm of Section 3.1. Finally, in section 4.3, we study the convexity of the set of local means $\mathcal{M}_h(P)$, which makes the relation between different definitions for the k -PDTM more transparent. Also, we clarify the connection between the k -PDTM and the Bregman-quantization method with $\psi_{P,h}$.

² $\psi_{P,h}$ is convex according to [CCSM11, Proposition 3.3].

4.1 The deterministic framework

A result from [CCSM11] states that the distance-to-measure can be expressed as a power function with centers in $\tilde{M}_h(P)$. This encourages us to reduce the number of centers to some $k \in \mathbb{N}^*$. In this section, we define the k -PDTM as the square root of the k -power function that is the best approximation from above of the square of the distance-to-measure in terms of the $L_1(P)$ norm. The convexity of the set $\tilde{M}_h(P)$ allows us to derive another equivalent definition for the k -PDTM. We prove that the centers associated to the optimal k -power approximation (that is, the k -PDTM) coincide with optimal codebooks for a k -means type method. They are actually to be found among the elements of the set of local means $\mathcal{M}_h(P)$.

4.1.1 The k -PDTM: a coreset for the DTM

The DTM defined as a power distance

In [CCSM11][Proposition 3.3], Chazal et al. provide an expression of the DTM in terms of sub-measures, given for every $x \in \mathbb{R}^d$ and $P_{x,h} \in \mathcal{P}_{x,h}$ by

$$d_{P,h}^2(x) = \inf_{Q \in \mathcal{P}_h(P)} Q \|x - u\|^2 = P_{x,h} \|x - u\|^2. \quad (4.1)$$

A *power distance* indexed on a set I is the square root of a *power function* $f_{\tau,\omega}$ defined on \mathbb{R}^d from a family of centres $\tau = (\tau_i)_{i \in I}$ and weights $\omega = (\omega_i)_{i \in I}$ by $f_{\tau,\omega} : x \mapsto \inf_{i \in I} \|x - \tau_i\|^2 + \omega_i^2$. A *k -power distance* is a power distance indexed on a finite set of cardinal $|I| = k$.

According to (4.1), the distance-to-measure can be written as a power distance. Namely, for every P in $\mathcal{P}_2(\mathbb{R}^d)$ and $x \in \mathbb{R}^d$,

$$d_{P,h}^2(x) = \inf_{Q \in \mathcal{P}_h(P)} \|x - m(Q)\|^2 + v(Q), \quad (4.2)$$

and the minimum is attained at any measure $Q = P_{x,h}$ in $\mathcal{P}_{x,h}(P)$. The centers are elements of $\tilde{M}_h(P)$.

As noted by Guibas et al. in [GMM11], this expression holds for the empirical DTM $d_{P_n,h}$. Then, $m(P_{n,x,h})$ corresponds to the barycentre of the $q = nh$ nearest-neighbours of x in \mathbb{X}_n , $\text{NN}_{q,\mathbb{X}_n}(x)$, and $v(P_{n,y,h}) = \frac{1}{q} \sum_{p \in \text{NN}_{q,\mathbb{X}_n}(x)} \|x - p\|^2$.

Introduction of the k -PDTM

The expression of the DTM as a power distance has already been exploited in [GMM11] and [BCOS15] to approximate it by n -power distances. In this paper, we approximate the distance-to-measure by a k -power distance whose graph lies above the graph of the DTM. With this aim in mind, we introduce a function $\omega_{P,h}$ for every τ in \mathbb{R}^d by

$$\omega_{P,h}(\tau) = \inf \left\{ \omega > 0 \mid \forall x \in \mathbb{R}^d, \|x - \tau\|^2 + \omega^2 \geq d_{P,h}^2(x) \right\}. \quad (4.3)$$

Morally, $\omega_{P,h}$ is the value at τ of the parabola that is closest to $d_{P,h}^2$ among all the parabolas centered at τ , above the graph of $d_{P,h}^2$. Then, any function defined for $\boldsymbol{\tau} = (\tau_1, \tau_2, \dots, \tau_k)$ by $x \mapsto \min_{i \in \llbracket 1, k \rrbracket} \|x - \tau_i\|^2 + \omega_{P,h}^2(\tau_i)$ is a k -power function that bounds $d_{P,h}^2$ above.

Equivalently, $\omega_{P,h}$ is defined by

$$\omega_{P,h}^2(\tau) = \sup_{x \in \mathbb{R}^d} d_{P,h}^2(x) - \|x - \tau\|^2. \quad (4.4)$$

The weight $\omega_{P,h}(\tau)$ is used as a penalization term of the Euclidean distance to a center τ in the k -means criterion. Note that $\omega_{P,h}(\tau)$ is large whenever τ is “far” from P in terms of the DTM to P . The set of optimal codebooks for this new penalized criterion is defined as follows.

Definition 4.3. The set $\text{OPT}(P, h, k)$ is defined by:

$$\text{OPT}(P, h, k) = \arg \min \left\{ P \min_{i \in [1, k]} \|u - \tau_i\|^2 + \omega_{P, h}^2(\tau_i) \mid \boldsymbol{\tau} = (\tau_1, \tau_2, \dots, \tau_k) \in \mathbb{R}^d \right\}^{(k)}.$$

We intend to define our approximation of the DTM as a function $x \mapsto \min_i \|x - \tau_i\|^2 + \omega_{P, h}^2(\tau_i)$, with $\boldsymbol{\tau} \in \text{OPT}(P, h, k)$. To this aim we need some structural results about $\text{OPT}(P, h, k)$. The following Lemma shows that $\text{OPT}(P, h, k)$ is included in $\tilde{\mathcal{M}}_h(P)^{(k)}$.

Lemma 4.4. *If $P \in \mathcal{P}_2(\mathbb{R}^d)$, then $\omega_{P, h}(\theta) < +\infty$ if and only if $\theta \in \tilde{\mathcal{M}}_h(P)$. Moreover, if $\theta \in \tilde{\mathcal{M}}_h(P)$, then there exists $Q \in \mathcal{P}_h(P)$ such that $\theta = m(Q)$ and $\omega_{P, h}^2(\theta) = v(Q)$.*

As a consequence, $\text{OPT}(P, h, k) \neq \emptyset$.

The proof of Lemma 4.4 is deferred to Section 4.4.1.

Definition 4.5. Let $P \in \mathcal{P}_2(\mathbb{R}^d)$, and $\boldsymbol{\tau} = (\tau_1, \dots, \tau_k)$ be in $\text{OPT}(P, h, k)$. The k -power distance-to-measure (k -PDTM) $d_{P, h, k}$ is defined by:

$$d_{P, h, k}^2(x) = \min_{i \in [1, k]} \|x - \tau_i\|^2 + \omega_{P, h}^2(\tau_i).$$

The k -PDTM at a point $x \in \mathbb{R}^d$ corresponds to a distance in \mathbb{R}^{d+1} of $(x, 0)$ to the compact set $\left\{ (\tau_i, \omega_{P, h}^2(\tau_i)) \mid \boldsymbol{\tau} \in \text{OPT}(P, h, k) \right\}$. In particular, the k -PDTM is automatically 1-Lipschitz.

Definition 4.5 introduces $d_{P, h, k}$ as the square root of a minimizer of the $L_1(P)$ norm $f \mapsto P|f - d_{P, h, k}^2|(u)$ among all the k -power functions f whose graph lies above the graph of the function $d_{P, h, k}^2$. Lemma 4.4 ensures that the optimal τ_i 's are to be found in $\tilde{\mathcal{M}}_h(\mathbb{R}^d)$. In practice, to find such an optimal $\boldsymbol{\tau}$, we have to optimize a cost function over all possible means of sub-measures of P with P -mass h , which is intractable. Thus, some further results on the k -PDTM are needed to ensure its tractability.

In Section 4.3, we prove that $\tilde{\mathcal{M}}_h(\mathbb{R}^d)$ coincides with the convex set $\overline{\mathcal{M}}_h(\mathbb{R}^d)$ whenever $P \in \mathcal{P}^K(\mathbb{R}^d)$ puts no mass on spheres nor on hyperplanes. Thus, for such measures, the optimal τ_i 's are to be found in the set of local means $\overline{\mathcal{M}}_h(\mathbb{R}^d)$. In the following section we prove that this property extends to all measures $P \in \mathcal{P}_2(\mathbb{R}^d)$.

Two equivalent definitions of the k -PDTM

In this section we prove that the optimal τ_i 's are to be found among the elements $m(P_{t_i, h})$ for t_i in \mathbb{R}^d and $P_{t_i, h} \in \mathcal{P}_{t_i, h}(P)$.

Definition 4.6. For all $P \in \mathcal{P}_2(\mathbb{R}^d)$ and $(P_i)_{i=1, \dots, k} \in \mathcal{P}_h(P)^{(k)}$, we define $R(P_1, \dots, P_k)$ by

$$R(P_1, \dots, P_k) = P \min_{i \in [1, k]} \|u - m(P_i)\|^2 + v(P_i).$$

With a slight abuse of notation, for every $\mathbf{t} \in \mathbb{R}^d \times \dots \times \mathbb{R}^d$, we denote the quantity $R(P_{t_1, h}, \dots, P_{t_k, h})$ by $R(\mathbf{t})$, where $P_{t_i, h} \in \mathcal{P}_{t_i, h}(P)$. A closely related notion to Definition 4.6 is the set of weighted Voronoi measures.

Definition 4.7. A set of *weighted Voronoi measures* associated to a distribution $P \in \mathcal{P}_2(\mathbb{R}^d)$, $(P_i)_{i=1, \dots, k} \in \mathcal{P}_h(P)^{(k)}$ and $h \in (0, 1]$ is a set $\left\{ \tilde{P}_{1, h}, \tilde{P}_{2, h}, \dots, \tilde{P}_{k, h} \right\}$ of $k \in \mathbb{N}^*$ non-negative sub-measures of P such that $\sum_{i=1}^k \tilde{P}_{i, h} = P$ and

$$\forall x \in \text{Supp}(\tilde{P}_{i, h}), \|x - m(P_i)\|^2 + v(P_i) \leq \|x - m(P_j)\|^2 + v(P_j), \forall j \in [1, k].$$

We denote by $\tilde{m}(\tilde{P}_{i,h}) = \frac{\tilde{P}_{i,h}u}{\tilde{P}_{i,h}(\mathbb{R}^d)}$ the expectation of $\tilde{P}_{i,h}$, with the convention that $\tilde{m}(\tilde{P}_{i,h}) = 0$ when $\tilde{P}_{i,h}(\mathbb{R}^d) = 0$.

Note that a set of weighted Voronoï measures can always be assigned to any $P \in \mathcal{P}_2(\mathbb{R}^d)$ and $(P_i)_{i=1,\dots,k} \in \mathcal{P}_h(P)^{(k)}$: it suffices to split \mathbb{R}^d into weighted Voronoï cells associated to the centers $(m(P_i))_{i \in \llbracket 1, k \rrbracket}$ and weights $(v(P_i))_{i \in \llbracket 1, k \rrbracket}$ ([BCY17, Section 4.4.2]) and split the remaining mass on the border of the cells in a measurable arbitrary way.

We will see in Section 4.3.3, that these weighted Voronoï measures actually coincide with the Bregman-Voronoï measures defined in Section 3.1 (see Definition 3.10), for the Bregman divergence associated with the function $\psi_{P,h}$. A key property of weighted Voronoï measures is the following.

Lemma 4.8. *Let $P \in \mathcal{P}_2(\mathbb{R}^d)$, and $(P_i)_{i=1,\dots,k} \in \mathcal{P}_h(P)^{(k)}$. Then*

$$R(m(\tilde{P}_{1,h}), \dots, m(\tilde{P}_{k,h})) \leq R(P_1, \dots, P_k),$$

with equality if and only if for all $i = 1, \dots, k$ such that $\tilde{P}_{i,h}(\mathbb{R}^d) \neq 0$, we have $P_i \in \mathcal{P}_{m(\tilde{P}_{i,h}),h}(P)$.

The proof of Lemma 4.8 is deferred to Section 4.4.1.

Theorem 4.9. *Let $P \in \mathcal{P}_2(\mathbb{R}^d)$, and $\tau \in \tilde{\mathcal{M}}_h(P)^{(k)}$. Then*

$$\tau \in \text{OPT}(P, h, k) \Leftrightarrow \forall i \in \llbracket 1, k \rrbracket, \quad \tau_i = m(P_{t_i,h}), \quad \omega_{P,h}^2(\tau_i) = v(P_{t_i,h}).$$

The codebook $\mathbf{t} = (t_1, t_2, \dots, t_k) \in \mathbb{R}^d$ in question is such that for all $(P_1, \dots, P_k) \in \mathcal{P}_h(P)^{(k)}$, we have $R(\mathbf{t}) \leq R(P_1, \dots, P_k)$.

Proof of Theorem 4.9. According to Lemma 4.4, we have that $\tau \in \text{OPT}(P, h, k)$ is equivalent to $\tau_i = m(P_i)$, $\omega_{P,h}^2(\tau_i) = v(P_i)$, and P_1, \dots, P_k such that $R(P_1, \dots, P_k) = \inf_{Q_1, \dots, Q_k} R(Q_1, \dots, Q_k)$. According to Lemma 4.8, for such a τ , denoting by $m_i = m(\tilde{P}_{i,h})$, we have $P_i \in \mathcal{P}_{m_i,h}(P)$. Hence $\omega_{P,h}^2(\tau_i) = v(P_{m_i,h})$ and $\tau_i = m(P_{m_i,h})$, for $P_{m_i,h} \in \mathcal{P}_{m_i,h}(P)$. Conversely, Lemma 4.8 ensures that if \mathbf{t} is an R -minimizer, then $\tau \in \text{OPT}(P, h, k)$, with $\tau_i = m(P_{t_i,h})$. \square

As a consequence, we may also define the k -PDTM via

$$d_{P,h,k}^2(x) = \min_{i \in \llbracket 1, k \rrbracket} \|x - m(P_{t_i,h})\|^2 + v(P_{t_i,h}), \tag{4.5}$$

where $\mathbf{t} \in \arg \min_{\mathbb{R}^d} R(\mathbf{t})$, and $P_{t_i,h} \in \mathcal{P}_{t_i,h}(P)$. Such an alternative definition is crucial for computing the k -PDTM. Indeed, if $P = P_n$ is the empirical distribution of an n -sample and $q = nh$ is an integer, then minimizing $R(\mathbf{t})$ can be carried out via a minimization over the set of barycenters of sets of q -nearest neighbours. We may also take advantage of Lemma 4.8 to design an iterative algorithm that converges toward a local minimum of R , as described in Section 4.2.

4.1.2 Proximity to the DTM

The complexity of a compact subset M of \mathbb{R}^d can be measured in terms of its ϵ -covering number, defined as the minimum number of balls in \mathbb{R}^d of size ϵ needed to cover M .

Here we show that the k -PDTM approximates the DTM provided that the covering number of M and the continuity modulus of the map $x \mapsto m(P_{x,h})$ are not too large.

Proposition 4.10. *Let $P \in \mathcal{P}^K(\mathbb{R}^d)$ for $K > 0$ and let $M \subset \mathcal{B}(0, K)$ be such that $P(M) = 1$. Let $f_M(\varepsilon)$ denote the ε covering number of M . Then we have*

$$0 \leq Pd_{P,h,k}^2(u) - d_{P,h}^2(u) \leq 2f_M^{-1}(k)\zeta_{P,h}(f_M^{-1}(k)), \quad \text{with} \quad f_M^{-1}(k) = \inf \{ \varepsilon > 0 \mid f_M(\varepsilon) \leq k \},$$

where $\zeta_{P,h}$ is the continuity modulus of $x \mapsto m(P_{x,h})$, that is

$$\zeta_{P,h}(\varepsilon) = \sup_{x,y \in M, \|x-y\| \leq \varepsilon} \sup_{P_{x,h} \in \mathcal{P}_{x,h}(P), P_{y,h} \in \mathcal{P}_{y,h}(P)} \{ |m(P_{x,h}) - m(P_{y,h})| \}.$$

Proof of Proposition 4.10. The first inequality comes from Proposition 4.18.

We now focus on the second bound. By definition of $d_{P,h,k}$, for all $x \in \mathbb{R}^d$ and $\mathbf{t} = (t_1, t_2, \dots, t_k) \in \mathbb{R}^{d(k)}$ we have $Pd_{P,h,k}^2(x) \leq P \min_{i \in [1,k]} \|u - m(P_{t_i,h})\|^2 + v(P_{t_i,h})$. Thus,

$$\begin{aligned} Pd_{P,h,k}^2(u) - d_{P,h}^2(u) &\leq P \min_{i \in [1,k]} \|u - m(P_{t_i,h})\|^2 + v(P_{t_i,h}) - d_{P,h}^2(u) \\ &= P \min_{i \in [1,k]} (d_{P,h}^2(t_i) - \|t_i\|^2) - (d_{P,h}^2(u) - \|u\|^2) + \langle u - t_i, -2m(P_{t_i,h}) \rangle \\ &\leq P \min_{i \in [1,k]} 2\langle u - t_i, m(P_{u,h}) - m(P_{t_i,h}) \rangle \\ &\leq 2P \min_{i \in [1,k]} \|u - t_i\| \|m(P_{u,h}) - m(P_{t_i,h})\|, \end{aligned}$$

where we used (4.8), Lemma 4.28 and the Cauchy-Schwarz inequality. Now choose t_1, \dots, t_k such that $M \subset \bigcup_{i \in [1,k]} \mathcal{B}(t_i, f_M^{-1}(k))$. The result follows. \square

When P is roughly uniform on its support, the quantities $f_M^{-1}(k)$ and $\zeta_{P,h}$ mostly depend on the dimension and radius of M . We focus on two cases in which Proposition 4.10 may be useful. First, the case where the distribution P has an ambient-dimensional support is investigated.

Corollary 4.11. *Assume that P has a density f satisfying $0 < f_{\min} \leq f \leq f_{\max}$. Then*

$$0 \leq Pd_{P,h,k}^2(u) - d_{P,h}^2(u) \leq C_{f_{\max}, K, d, h} k^{-2/d}.$$

The proof of Corollary 4.11 is given in Section 4.4.1. Note that no assumptions on the geometric regularity of M are required for Corollary 4.11 to hold. In the case where M has a lower-dimensional structure, more regularity is required, as for instance in the following corollary.

Corollary 4.12. *Suppose that P is supported on a compact d' -dimensional \mathcal{C}^2 -submanifold of $\mathcal{B}(0, K)$, denoted by N . Assume that N has positive reach ρ , and that P has a density $0 < f_{\min} \leq f \leq f_{\max}$ with respect to the volume measure on N . Moreover, suppose that P satisfies, for all $x \in N$ and positive r ,*

$$P(\mathcal{B}(x, r)) \geq cf_{\min} r^{d'} \wedge 1. \quad (4.6)$$

Then, for $k \geq c_{N, f_{\min}}$ and $h \leq c_{N, f_{\min}}$, we have $0 \leq Pd_{P,h,k}^2(u) - d_{P,h}^2(u) \leq C_{N, f_{\min}, f_{\max}} k^{-2/d'}$.

Note that (4.6), also known as (cf_{\min}, d') -standard assumption, is usual in set estimation (see, e.g., [CGLM15]). In the submanifold case, it may be thought of as a condition preventing the boundary from being arbitrarily narrow. This assumption is satisfied for instance in the case where ∂N is empty or is a \mathcal{C}^2 $d' - 1$ -dimensional submanifold (see, e.g., [AC17, Corollary 1]). An important feature of Corollary 4.12 is that this approximation bound does not depend on the ambient dimension. The proof of Corollary 4.12 may be found in Section 4.4.1.

4.1.3 Wasserstein stability for the k -PDTM

Next we assess that our k -PDTM shares with the DTM the key property of robustness to noise.

Proposition 4.13. *Let $P \in \mathcal{P}^K(\mathbb{R}^d)$ for some $K > 0$ and $Q \in \mathcal{P}_2(\mathbb{R}^d)$. Set $d_{Q,h,k}^2$ a k -PDTM for Q . Then, $P \left| d_{Q,h,k}^2(u) - d_{P,h}^2(u) \right|$ is bounded above by $B_{P,Q,h,k}$ with*

$$B_{P,Q,h,k} = 3\|d_{Q,h}^2 - d_{P,h}^2\|_{\infty, \mathcal{B}(0,K)} + P d_{P,h,k}^2(u) - d_{P,h}^2(u) + 4W_1(P, Q) \sup_{s \in \mathbb{R}^d} \|m(P_{s,h})\|.$$

Note that Lemma 4.36 provides a bound on $\|m(Q_{s,h})\|$ whenever Q is sub-Gaussian. Moreover, the term $\|d_{Q,h}^2 - d_{P,h}^2\|_{\infty, \mathcal{B}(0,K)}$ can be bounded above by $\|d_{Q,h} - d_{P,h}\|_{\infty, \mathcal{B}(0,K)} \|d_{Q,h} + d_{P,h}\|_{\infty, \mathcal{B}(0,K)}$. From the stability property of the DTM, this term gets bounded above by $W_2(P, Q)$, up to a constant dependent on K . Also, upper-bounds for the deviation of the k -PDTM to the DTM associated to P have been derived in the previous subsection.

The proof of Proposition 4.13 can be found in Section 4.4.1.

4.1.4 Geometric inference with the k -PDTM

As detailed in [CCSM11, Section 4], under suitable assumptions, the sublevel sets of the distance-to-measure are close enough to the sublevel sets of the distance to its support. Thus they allow us to infer the geometric structure of the support. As stated below, this is also the case when replacing the distance-to-measure with the k -PDTM.

Proposition 4.14. *Let M be a compact set in $\mathcal{B}(0, K)$ such that $P(M) = 1$. Moreover, assume that there exists d' such that, for every $p \in M$ and $r \geq 0$,*

$$P(\mathcal{B}(p, r)) \geq C(P)r^{d'} \wedge 1. \tag{4.7}$$

Let Q be a Borel probability measure (thought of as a perturbation of P), and let Δ_P^2 denote $P d_{Q,h,k}^2(u)$. Then, we have

$$\sup_{x \in \mathbb{R}^d} |d_{Q,h,k}(x) - d_M(x)| \leq \max \left\{ C(P)^{-\frac{1}{d'+2}} \Delta_P^{\frac{2}{d'+2}}, 2\Delta_P, W_2(P, Q)h^{-\frac{1}{2}} \right\},$$

where W_2 denotes the Wasserstein distance.

The proof of Proposition 4.14 can be found in Section 4.4.1.

According to Theorem 1.4, in this context, the persistence diagram associated with the sublevel sets of the k -PDTM $d_{Q,h,k}$ is close to the persistence diagram associated with the sublevel sets of the distance to M in terms of the bottleneck distance. Recall that the sublevel sets of the distance to M coincides with its offsets. Thus, it is possible to infer the persistent homology of M from the k -PDTM.

According to [CCSM11, Corollary 4.8], if P satisfies (4.7), then

$$\|d_{Q,h} - d_M\|_{\infty} \leq \left(\frac{h}{C(P)} \right)^{\frac{1}{d'}} + W_2(P, Q)h^{-\frac{1}{2}}.$$

Hence, Proposition 4.14 ensures that the k -PDTM achieves roughly the same performance as the distance-to-measure provided that $d_{Q,h,k}^2$ is small enough on the support M to be inferred. As will be shown in the following section, this will be the case if Q is an empirical measure drawn close to the targeted support.

4.2 The statistical framework

In this section, $P \in \mathcal{P}^K(\mathbb{R}^d)$ is a distribution supported on a compact set \mathcal{K} to be inferred. We have at our disposal an n -sample $\mathbb{X}_n = \{X_1, X_2, \dots, X_n\}$ from a modification $Q \in \mathcal{P}_2(\mathbb{R}^d)$ of P .

An approximation of the k -PDTM $d_{Q,h,k}$, is given by the *empirical k -PDTM* $d_{Q_n,h,k}$, where $Q_n = \sum_{i=1}^n \frac{1}{n} \delta_{X_i}$ is the empirical measure from \mathbb{X}_n . Note that when $k = n$, the empirical k -PDTM $d_{Q_n,h,n}$ coincides with the q -witnessed distance $d_{\mathbb{X}_n,q}^w$ introduced in [GMM11] (for q defined by $q = nh$). This approximation of the DTM is defined for every $x \in \mathbb{R}^d$ by

$$(d_{\mathbb{X}_n,q}^w)^2(x) = \min_{i \in [1,n]} \|x - m(Q_n X_i, h)\|^2 + v(Q_n X_i, h).$$

The equality of these two notions comes from the fact that for every $\mathbf{t} = (t_1, t_2, \dots, t_n) \in \mathbb{R}^{d(n)}$, we have:

$$\begin{aligned} Q_n \min_{i \in [1,n]} \|u - m(Q_n t_i, h)\|^2 + v(Q_n t_i, h) &\geq Q_n \|u - m(Q_n u, h)\|^2 + v(Q_n u, h) \\ &= Q_n (d_{\mathbb{X}_n,q}^w)^2(u), \end{aligned}$$

and that equality holds for $\mathbf{t} = (X_1, X_2, \dots, X_n)$.

In the following, we investigate the quality of approximation of the DTM $d_{P,h}$ with the empirical k -PDTM $d_{Q_n,h,k}$, when Q is defined as the convolution of P with a sub-Gaussian distribution with variance σ^2 . Within this context, according to Lemma 4.37, Q is sub-Gaussian with variance $V^2 = (K + \sigma)^2$. In particular, we recover k -means type rates (i.e. of order $\sqrt{\frac{kd}{n}}$ up to a logarithmic factor) for the approximation of the k -PDTM with its empirical counterpart in $L_1(Q)$. Also, we derive a Lloyd-type algorithm to compute a local minimizer for the empirical cost function associated with the empirical k -PDTM. We conclude this section with numerical experiments.

4.2.1 An algorithm for the empirical k -PDTM

The computation of an optimal codebook for the k -means method is NP-hard, thus intractable in practice. Likewise, computing a minimizer of our empirical cost function

$$(t_1, t_2, \dots, t_k) \mapsto Q_n \min_{i \in [1,k]} \|x - m(Q_n t_i, h)\|^2 + v(Q_n t_i, h)$$

is intractable. In this section, we display a Lloyd-type algorithm to compute a local minimum of this empirical criterion. From such a local minimizer, we derive an approximation of the k -PDTM.

Recall that Lloyd's algorithm consists, from a set of k centres, in splitting the data \mathbb{X}_n according to the Voronoi diagram associated to these centres. The second step consists in replacing the centres by, for every centre, the mean of the data within the associated Voronoi cell.

The new version that we develop is based on the local means $m(Q_n t, h)$ (denoted by $c(t)$) and variances $v(Q_n t, h)$ (denoted by $\omega^2(t)$) at points $t \in \mathbb{R}^d$. We assume that the mass parameter $q = hn$ is an integer. Then, for every $t \in \mathbb{R}^d$, $c(t) = \frac{1}{q} \sum_{i=1}^q X_i(t)$, where $X_i(t)$ is one of the i -th nearest neighbour of t in \mathbb{X}_n and $\omega^2(t) = \frac{1}{q} \sum_{i=1}^q (X_i(t) - c(t))^2$. As a consequence, the computation of $c(t)$ and $\omega^2(t)$ boils down to a q -nearest neighbors query, which is tractable.

In our new version of the Lloyd algorithm, weights are associated to the centres t_i s. Actually, we do not compute the Voronoï cell associated with the centres t_i in \mathbb{R}^d , but the Voronoï cells associated with the centres $(c(t_i), \omega(t_i))$ in \mathbb{R}^{d+1} .

The weighted Voronoï cells we consider are denoted by $\mathcal{C}(t_i)$ and defined by:

$$\mathcal{C}(t_i) = \{x \in \mathbb{X}_n \mid \|x - c(t_i)\|^2 + \omega^2(t_i) \leq \|x - c(t_l)\|^2 + \omega^2(t_l) \forall l \neq i\} \cap \bigcap_{j < i} \mathcal{C}(t_j)^c.$$

We use the notation $|\mathcal{C}(t)|$ for the cardinal number of $\mathcal{C}(t)$.

Algorithm 4.1: Local minimum algorithm

Input : \mathbb{X}_n an n -sample from Q , q and k ;
Initialization
 Sample t_1, t_2, \dots, t_k from \mathbb{X}_n without replacement. ;
while the t_i s vary make the following two steps :
Decomposition into weighted Voronoi cells.
for j in 1..n:
 Add X_j to the $\mathcal{C}(t_i)$ (for i as small as possible) satisfying
 $\|X_j - c(t_i)\|^2 + \omega^2(t_i) \leq \|X_j - c(t_l)\|^2 + \omega^2(t_l) \forall l \neq i$;
Computation of the new centers and weights.
for i in 1..k:
 $t_i = \frac{1}{|\mathcal{C}(t_i)|} \sum_{X \in \mathcal{C}(t_i)} X$;
Output : (t_1, t_2, \dots, t_k)

As aforementioned, the algorithm furnishes an approximation of the k -PDTM. Since the algorithm does not converge to the optimal centers, running the algorithm several times and storing the best solution in terms of the empirical cost function will be well suited.

Note that such an algorithm coincides with the Bregman clustering algorithm proposed by Banerjee et al. in [BMDG05], for the Bregman divergence associated with the convex function $\psi_{Q_n, h}$.

Proposition 4.15. *This algorithm converges to a local minimum of*

$$\mathbf{t} \mapsto Q_n \min_{i \in [1, k]} \|u - m(Q_n t_i, h)\|^2 + v(Q_n t_i, h).$$

The proof of Proposition 4.15 can be found in Section 4.4.2.

4.2.2 Proximity between the k -PDTM and its empirical version

In [CCSM11], the difference between the DTM and the empirical DTM in terms of the L_∞ norm is upper-bounded by $\frac{1}{\sqrt{h}} W_2(Q, Q_n)$. According to [FG15], its convergence to 0 occurs at a rate $n^{-\frac{1}{d}}$. Some improvements are obtained in [CMM16] for the approximation of the DTM with its empirical counterpart, in terms of the L_∞ -norm. The authors get the rate $n^{-\frac{1}{2}}$. The dependence on h is related to the continuity modulus of the quantile function of $\|x - X\|^2$ for $X \sim Q$. This underlines the fact that the empirical DTM is a good approximation of the DTM.

In the following, we prove that the empirical squared k -PDTM is also a good approximation of the squared k -PDTM, in terms of the $L_1(P)$ norm. Without further assumptions, the rate attained is $\frac{kd}{\sqrt{n}}$ up to a logarithmic term. It corresponds to the rate for the method of k -means derived in [BDL08]. The additional term in $\frac{\sigma}{\sqrt{h}}$ comes from the bias introduced by considering the error in terms of the $L_1(P)$ -norm, in place of $L_1(Q)$.

Theorem 4.16. *Let P be supported on $M \subset \mathcal{B}(0, K)$. Assume that we observe X_1, \dots, X_n such that $X_i = Y_i + Z_i$, where the Y_i 's and Z_i 's are all independent, Y_i is sampled from P and Z_i is sub-Gaussian with variance σ^2 , with $\sigma \leq K$. Let Q_n denote the empirical distribution associated with the X_i 's. Then, for any $p > 0$, with probability larger than $1 - 10n^{-p}$, we have*

$$|P(d_{Q_n, h, k}^2(u) - d_{Q, h, k}^2(u))| \leq C\sqrt{kd} \frac{K^2((p+1)\log(n))^{\frac{3}{2}}}{h\sqrt{n}} + C\frac{K\sigma}{\sqrt{h}}.$$

A proof of Theorem 4.16 is given in Section 4.4.2. Theorem 4.16, combined with Proposition 4.13, allows us to choose k in order to minimize $|Pd_{Q_n, h, k}^2(u) - d_{P, h}^2(u)|$. Indeed, in the framework of Corollaries 4.11 and 4.12 where the support has intrinsic dimension d' , such a minimization boils down to optimizing a quantity of the form

$$\frac{C\sqrt{k}K^2((p+1)\log(n))^{\frac{3}{2}}}{h\sqrt{n}} + C_{P, h}k^{-\frac{2}{d'}},$$

with C a positive absolute constant and $C_{P, h}$ a positive constant which depends on P and h .

Hence the choice $k \sim n^{\frac{d'}{d'+4}}$ ensures that for n large enough, only $n^{\frac{d'}{d'+4}}$ points are sufficient to approximate well the sub-level sets of the distance to support. For surface inference ($d' = 2$), this amounts to computing the distance to $n^{\frac{1}{3}}$ points rather than n , which might save some time. Note that when d' is large, smaller choices of k , though suboptimal for our bounds, would nonetheless give the right topology for large n . In some sense, Theorem 4.16 shows that there is a k above which no increase of precision can be expected.

More precisely,

Proposition 4.17. *With the same setting as Theorem 4.16, if M is a submanifold with intrinsic dimension $d' \geq 1$, then:*

$$|Pd_{Q_n, h, k}^2(u) - d_{P, h}^2(u)| \leq C\sqrt{kd} \frac{K^2((p+1)\log(n))^{\frac{3}{2}}}{h\sqrt{n}} + C\frac{K\sigma}{h} + C_{P, h}k^{-\frac{2}{d'}}.$$

Moreover, with the optimal choice $k \sim n^{\frac{d'}{d'+4}}$ we have

$$|Pd_{Q_n, h, k}^2(u) - d_{P, h}^2(u)| \leq C_{P, h}\sqrt{dn}^{-\frac{2}{d'+4}} \frac{K^2((p+1)\log(n))^{\frac{3}{2}}}{h} + C\frac{K\sigma}{h}.$$

The proof of Proposition 4.17 is to be found in Section 4.4.2.

4.2.3 Numerical illustrations

Topological inference from the sub-level sets of the k -PDTM

The function k -PDTM has been introduced to approximate the DTM associated to a distribution P , in order to recover the topology of its support $\mathcal{K} = \text{Supp}(P)$. This topological information is contained in the sub-level sets of the k -PDTM. We compare our approximation of the DTM with the k -witnessed distance, introduced in [GMM11]. With this aim in mind, we consider the same example, a 2-dimensional ∞ -shaped compact set that they call “sideways”.

Set P to be the uniform distribution on this sideways (defined as the “contour” of two spheres, one of radius $\sqrt{2}$ and the other one of radius $\sqrt{\frac{9}{8}}$). Then, Q is defined by the convolution of P with a Gaussian distribution $\mathcal{N}(0, \sigma^2)$, for $\sigma = 0.45$. We generate a $n = 6000$ -sample from Q . We plot in grey the r -sub-level set of the q -witnessed distance and in purple the r -sub-level set of an




Figure 4.1: Sideways - 6000-sample




Figure 4.2: $k = 100$




Figure 4.3: $k = 300$

approximation of $d_{Q_n, h, k}$ for $r = 0.24$, $q = 50$ nearest-neighbors and $h = \frac{q}{n}$. The approximation of $d_{Q_n, h, k}$ is obtained after running our algorithm 10 times, with at most 10 iterations. We keep the best codebook for the empirical cost function.

Choosing $k = 100$ points leads to a too sparse approximation of the sub-level sets of the DTM. When $k = 300$, there are small holes in the r -sub-level set of the k -PDTM. But these holes will disappear quickly when the radius r gets larger, before the two holes get filled. Such a sub-level set will be made of one connected component and two holes. So, for some well-chosen parameter k , it is possible to infer the ∞ -shaped topology of the compact set \mathcal{K} .

A new quantization trimming method to deal with outliers

The second major interest of our method consists in providing a new version of the k -means method, adapted to quantization, in presence of outliers or clutter noise.

In this section, we consider Q , the convolution of P with an isotropic centered Gaussian distribution in \mathbb{R}^3 . We generate 200 points according to Q and 100 additional points of clutter noise, so that an $n = 300$ -sample is available. We set $q = 20$ (that is $h = \frac{20}{300}$) and $k = 30$, the initial number of centers for all of the Lloyd-type algorithms we will consider.

Failure of k -means and trimmed k -means. In this context, neither the k -means method (with an additional trimming step), nor the trimmed k -means method of Cuesta et al. in [CAGM97] perform in the detection of outliers. In order to illustrate this failure, we compute an optimal codebook for k -means, via the Lloyd algorithm. We sort the values $l_j = \min_{i \in \llbracket 1, k \rrbracket} \|X_j - c_i^*\|$ in ascending order, for $j \in \llbracket 1, n \rrbracket$. Roughly, we want to say that a point X_j is an outlier when l_j is large, otherwise, we want to say that X_j is a signal point. In Figure 4.4, we sort all of the l_j s in the increasing order. Then, we plot their cumulative sums. As we can see, there is no breaking point in the curve. It means that the k -means criterion does not allow a clear distinction between signal and noise. When we know the true amount of signal (here 200 points) and decide not to assign the $n - 200$ points for which the l_j are the largest to clusters, we get Figure 4.5. Manifestly, it does not work.

The trimmed k -means method of [CAGM97] is especially designed to face the problem of outliers in the context of clustering. Unfortunately, in the context of quantization (k large), the method fails. For different trimming parameters h' , we have implemented the trimmed k -means algorithm (Algorithm 3.1 for the Bregman divergence equal to the square of the Euclidean norm). We have computed the trimmed k -means criterion associated with the codebook returned by the algorithm, for different trimming parameters h' . In Figure 4.6, we have plotted this as a function of the number of points assigned to clusters, nh' . Again, there is no breaking point in the curve. The quantization associated with the trimming parameter colored in blue in Figure 4.6 is available in Figure 4.7. Again, it fails drastically.

Note that the selection of the trimming parameter for the trimmed k -means method is more expensive in terms of computational time than for the first method. Indeed, the first method requires to compute only one optimal codebook (the optimal codebook for the cost function $\mathbf{c} \rightarrow Q_n \min_{i \in \llbracket 1, k \rrbracket} \|u - c_i\|^2$). Then, we trim by computing the average of the nh' smallest values




Figure 4.4: Mean cost




Figure 4.5: Clustering with k -means

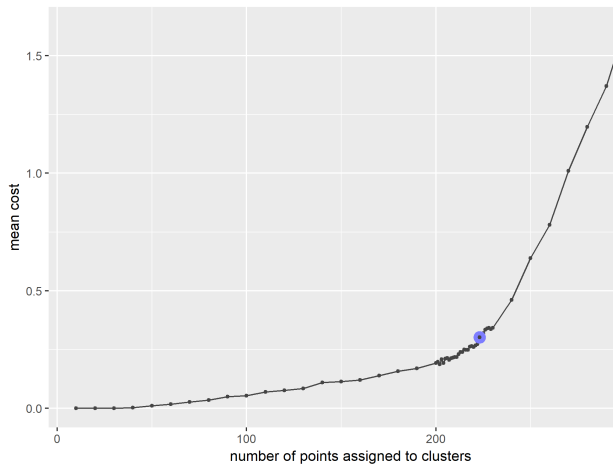


Figure 4.6: Mean cost




Figure 4.7: Trimmed clustering

of $\min_{i \in [1, k]} \|X_j - c_i\|$ among the points X_j of the sample. On the other hand, the second method requires to compute as many optimal codebooks as trimming parameters (the optimal codebooks for the cost functions $\mathbf{c} \rightarrow \inf_{Q_{n,h}} Q_{n,h} \min_{i \in [1, k]} \|u - c_i\|^2$, for $Q_{n,h} \in \mathcal{P}_h(Q_n)$). Thus, it requires to run the algorithm many times, while the first method requires to run it only once.

Quantization with the k -PDTM. The methods of k -means and trimmed k -means fail drastically in the context of quantization, for data corrupted with outliers. In this part, we present numerical illustrations obtained with our new quantization method based on the cost function $R : \mathbf{c} \mapsto Q_n \min_{i \in [1, k]} \|u - c_i\|^2 + \omega_{Q_n, h}^2(c_i)$, or equivalently, on the cost function $\mathbf{t} \rightarrow Q_n \min_{i \in [1, k]} \|u - m(Q_{n, t_i, h})\|^2 + v(Q_{n, t_i, h})$.

The parameter h for the k -PDTM is set to $q = 20$ nearest neighbors, that is $h = \frac{20}{300}$. Again, there are 200 points of signal and 100 points of clutter noise. The initial number of clusters for the algorithms is set to $k = 30$. Also, the maximum number of iterations is set to 100, and the best codebook is kept among 10 trials.

Similarly to the previous part, when we want to trim, there are two methods to be compared. The first method consists in computing an optimal codebook for the cost function R , and to trim

after. For this method, the cost is plotted in Figure 4.8 as a function of the number of points $h'n$ preserved in the sample. We select the parameter $h'n = 201$. The quantization associated to this parameter, with the k -PDTM is available in Figure 4.9. Unlike the previous part, there is a breakdown point in the curve representing the cost as a function of the trimming parameter. The quantization method associated with this trimming parameter is effective in the detection of outliers or clutter noise.




Figure 4.8: Mean cost




Figure 4.9: Quantization with the k -PDTM

The second method consists in adapting the trimmed k -means method to our cost function. That is to say, we implement the trimmed clustering method developed in the previous Chapter, in Section 3.1, to the Bregman divergence associated to the convex function $\psi_{P,h}$. It consists in minimizing the cost function $c \mapsto \inf_{Q_{n,h}} Q_{n,h} \min_{i \in \llbracket 1, k \rrbracket} \|u - c_i\|^2 + \omega_{Q_{n,h}}^2(c_i)$ where $Q_{n,h} \in \mathcal{P}_h(Q_n)$. The cost is represented in Figure 4.10 as a function of the trimming parameter $h'n$. The corresponding optimal quantization is available in Figure 4.11. Again, this method requires to compute as many optimal codebooks as the amount of trimming parameters desired. An advantage of this method is that there are more centers preserved than for the method consisting in trimming the k -PDTM once computed, in Figure 4.9.




Figure 4.10: Mean cost




Figure 4.11: Trimmed k -PDTM quantization

In order to compare the different methods, we use the normalized mutual information of

[SJ02]. Its definition was recalled in Section 3.1.5; see Definition 3.22.

The normalized mutual information is computed between two clusterings of type “signal vs noise”. The first clustering is associated with the true labels. The second one is associated to the labels obtained after any other outliers detection process. The label 0 is assigned to points considered as clutter noise and the label 1 is assigned to points considered as signal. The normalized mutual information is high whenever the two clusterings are close. At most, the NMI equals to 1. It corresponds to the case for which the two labelings coincide.

The NMI computed for the four aforementioned methods is available in the following table, in Figure 4.12. It confirms that unlike k -means, our method with the k -PDTM is perfectly adapted to separate signal from noise, in the context of quantization.

k -means	trimmed k -means	k -PDTM	trimmed k -PDTM
0.663	0.649	0.969	0.969

Figure 4.12: Normalized mutual information : “signal vs noise”

Topological inference from data corrupted with outliers or clutter noise

We proved that the k -PDTM was particularly well adapted to topological inference for data generated close to the compact set \mathcal{K} . Indeed, its sub-level sets help in recovering the topology. Its trimmed version is well-suited to detect outliers or clutter noise. In the following, we consider these two approaches at the same time. The performance of our method is illustrated through persistence diagrams. The quality of approximation to the persistence diagram associated to the true shape (the sideways) is measured in terms of the bottleneck distance.

The persistence diagram associated to the distance to the sideways should contain three points. One of them is a red point corresponding to the connected component whose birth time is 0 and death time is ∞ . Indeed, the sideways contains a single connected component. Furthermore, the diagram contains two green points corresponding to the two holes in the sideways. They should both be born at time 0, one should die at time $\sqrt{2} \simeq 1.414214$ and the other one should die at time $\sqrt{\frac{9}{8}} \simeq 1.06066$.

In Figure 4.13, we have plotted the persistence diagram associated with the sub-level sets of the k -PDTM. In Figure 4.14, we have represented the persistence diagram associated with the sub-level sets of the k -PDTM, after thresholding (c.f. the first method of quantization with the k -PDTM). Roughly, we kept the centers and weights associated with points considered as signal in our quantization method with the k -PDTM. Then, in Figure 4.15, we have plotted the persistence diagram associated with the sub-level sets of the trimmed k -PDTM. We recover the three aforementioned main components : the connected component and the two holes. The additional points are close to the diagonal, thus should be considered as noise due to the fact that the k -PDTM is a discrete approximation of the distance to the original compact set.

Now, we replace the k -PDTM with the q witnessed distance. We compare our method to the corresponding method. In Figure 4.16, we have plotted the persistence diagram associated to the k -witnessed distance of [CAGM97], whereas in Figure 4.17, we plotted the persistence diagram associated to its sub-level sets, after thresholding.

The NMI associated with the k -witnessed distance after thresholding equals to 0.9687825. Its performance is roughly the same as the performance of k -PDTM for $k = 40$, in terms of outlier detection. Recall that the k -witnessed function coincides with the n -PDTM, for n the




Figure 4.13: Persistence diagram for the k -PDTM without thresholding




Figure 4.14: Persistence diagram for the k -PDTM after thresholding




Figure 4.15: Persistence diagram for the trimmed k -PDTM




Figure 4.16: Persistence diagram – Witnessed distance




Figure 4.17: Persistence diagram – Witnessed distance after thresholding

sample size. The main advantage of using the k -PDTM for k much smaller than n is in terms of computational time and storage capacity for the topological descriptors such as the persistence diagrams. Indeed, the persistence diagrams contain fewer points but still contain the relevant features (although a little bit shifted with respect to the points in the persistence diagram associated to the original shape).

Some discussion on the selection of h

The k -PDTM is built from the DTM, which depends on a mass parameter $h \in [0, 1]$. In the discrete case, the number of nearest neighbors to consider is given by $q = nh$. Consequently, there is a large family of k -PDTM functions or equivalently of quantization procedures that can be computed from a dataset. It is pertinent to wonder if there are some values h for which our method performs better. Such a question is tough. Indeed, even for the DTM, the question of the selection of h has not been tackled yet, though tough.

Over a first phase, we compute the normalized mutual information (see Definition 3.22) of the clustering “signal vs noise” in Figure 4.18 for different values of q . There is a window of mass parameters h for which the NMI equals 1, meaning that we recover the true decomposition “signal vs noise”. For very small values h , our cost function is very close to the one associated to the k -means method. As illustrated above, such a method does not perform well when dealing with outliers or clutter noise. On the contrary, when q is large, the DTM at a point x measures its distance to the mean of the distribution Q_n . However, points will be considered as outliers when




Figure 4.18: Normalised mutual information

they lie outside a ball of Q_n -mass h' . Thus, for large values of h , the method fails drastically to detect outliers. The mass parameter h should be small, but not too small.

In Figure 4.27, we have plotted the quantization we obtained using the trimmed k -means for different values q . Some associated persistence diagrams are available in Figure 4.32. The proximity between the diagrams associated to our method and the diagram associated to the underlying compact set \mathcal{K} is available in Figure 4.33. Again, the mass parameter h should be small, but not too small, to recover the topology of \mathcal{K} . There is still work to do for the selection of the best h , from data.




Figure 4.19: $q = 20$




Figure 4.20: $q = 50$




Figure 4.21: $q = 100$




Figure 4.22: $q = 140$




Figure 4.23: $q = 160$




Figure 4.24: $q = 200$




Figure 4.25: $q = 290$




Figure 4.26: $q = 300$

Figure 4.27: Quantization with the trimmed k -PDTM for different mass-parameters q




Figure 4.28: $q = 20$




Figure 4.29: $q = 50$




Figure 4.30: $q = 100$




Figure 4.31: $q = 200$

Figure 4.32: Persistence diagrams for the trimmed k -PDTM for different mass-parameter $q = hn$

q	homology of type 0	homology of type 1	homology of type 2
20	0.284	0.640263	0.99962
50	0.755053	0.707107	0
10	1.78225	0.707107	0
200	5.02832	0.707107	0

Figure 4.33: Bottleneck distance to the true diagram

4.3 The set of local means – A strong connection with the distance-to-measure

In [CCSM11, Proposition 3.3], Chazal et al. provided the following expression for the DTM in terms of local means: for all $P \in \mathcal{P}(\mathbb{R}^d)$, $x \in \mathbb{R}^d$ and $P_{x,h} \in \mathcal{P}_{x,h}(P)$,

$$d_{P,h}^2(x) = P_{x,h} \|x - u\|^2 = \|x - m(P_{x,h})\|^2 + v(P_{x,h}). \tag{4.8}$$

The measures in $\mathcal{P}_{x,h}(P)$ are the minimizers of the map $P' \mapsto P' \|x - u\|^2$ among all distributions P' such that hP' is a sub-measure of P with P -mass h . So, another characterization of the DTM is as follows.

Proposition 4.18 ([CCSM11, Proposition 3.3]). *If $P \in \mathcal{P}_2(\mathbb{R}^d)$, then for all $x \in \mathbb{R}^d$, we have:*

$$d_{P,h}^2(x) = \inf_{y \in \overline{\mathbb{R}^d}} \inf_{P_{y,h} \in \mathcal{P}_{y,h}(P)} \|x - m(P_{y,h})\|^2 + v(P_{y,h})^3,$$

and the minimum is attained at $y = x$ and any measure $P_{x,h} \in \mathcal{P}_{x,h}(P)$.

In this way, the distance-to-measure is strongly related to the set of means of sub-measures of P , and more precisely, to the set of local means. So too is the quantization method that we introduced in this Chapter.

In Section 4.1, the quantization method we proposed (or equivalently, the k -PDTM) was defined from minimizers of the following cost function:

$$R : \mathbf{c} \mapsto P \min_{i \in \llbracket 1, k \rrbracket} \|u - c_i\|^2 + \omega_{P,h}^2(c_i).$$

According to Lemma 4.4, the set of points c that satisfy $\omega_{P,h}^2(c) < \infty$ coincides with the convex set $\tilde{\mathcal{M}}_h(P)$. As a consequence, optimal codebooks are to be found in $\tilde{\mathcal{M}}_h(P)$, the set of means of sub-measures of P with P -mass h .

In this section, we derive an expression for $\tilde{\mathcal{M}}_h(P)$ and for the convex hull of $\overline{\mathcal{M}}_h(P)$ (the “closure” of the set of local means), as an intersection of half spaces associated with extremal points. The extremal points are defined as means of sub-measures of P , coinciding with P on half-spaces of P -mass h . Morally, they correspond to the limit of a local mean when the center of the ball goes to ∞ in one direction. These expressions for $\tilde{\mathcal{M}}_h(P)$ and $\text{Conv}(\overline{\mathcal{M}}_h(P))$ coincide. As a consequence, whenever $\overline{\mathcal{M}}_h(P)$ is convex, $\tilde{\mathcal{M}}_h(P)$ and $\overline{\mathcal{M}}_h(P)$ will coincide. Then, automatically, any minimizer of R will belong to the “closure” of the set of local means $\overline{\mathcal{M}}_h(P)$.

³The set $\overline{\mathbb{R}^d}$ is a closure of \mathbb{R}^d . For $y \in \overline{\mathbb{R}^d} \setminus \mathbb{R}^d$, the measure $hP_{y,h}$ is the restriction of P on an half-space of P -mass h , as defined in Section 4.3.1.

Over a second phase, under different settings, we prove that the set $\overline{\mathcal{M}}_h(P)$ is convex. Last but not least, we showcase the relation between the construction of the k -PDTM and the Bregman quantization of [BMDG05], with the Bregman divergence associated with $\psi_{P,h}$.

4.3.1 Local means – From balls to half-spaces

By making the center of a ball go to ∞ along a direction $v \in S(0, 1)$ (the unit sphere of \mathbb{R}^d) such that the ball keeps a fixed mass h , we obtain sub-measures of P supported on half-spaces.

For $v \in S(0, 1)$, we denote by v_∞ the infinite point associated to the direction v . It can be seen as a limit point $\lim_{\lambda \rightarrow +\infty} \lambda v$. Then, we denote $\overline{\mathbb{R}}^d = \mathbb{R}^d \cup \{v_\infty \mid v \in S(0, 1)\}$. Note that we can equip $\overline{\mathbb{R}}^d$ with the metric $d_{\overline{\mathbb{R}}^d}$ defined by

$$d_{\overline{\mathbb{R}}^d}(x, y) = \|\phi(x) - \phi(y)\|_2,$$

with $\phi(x) = \frac{x}{\sqrt{1+\|x\|_2^2}}$ when $x \in \mathbb{R}^d$ and $\phi(v_\infty) = v$ for all $v \in S(0, 1)$.

For this metric, a sequence $(x_n)_{n \in \mathbb{N}}$ of $\overline{\mathbb{R}}^d$ converges to v_∞ if and only if $\lim_{n \rightarrow +\infty} \|x_n\|_2 = +\infty$ and $\lim_{n \rightarrow +\infty} \frac{x_n}{\|x_n\|_2} = v$ with the convention $\frac{w_\infty}{\|w_\infty\|_2} = w$ for all $w \in S(0, 1)$.

Indeed, provided that $\lim_{n \rightarrow +\infty} \|x_n\|_2 = +\infty$,

$$\lim_{n \rightarrow +\infty} \left\| \phi(x_n) - \frac{x_n}{\|x_n\|_2} \right\|_2 = \lim_{n \rightarrow +\infty} \left| \frac{\|x_n\|_2}{\sqrt{1 + \|x_n\|_2^2}} - 1 \right| = 0.$$

For $v \in S(0, 1)$, we can define a notion analogous to the pseudo-distance for half-spaces $c_{P,h}(v)$ by

$$c_{P,h}(v) = \sup \left\{ c \in \mathbb{R} \mid P \left(\left\{ x \in \mathbb{R}^d \mid \langle x, v \rangle > c \right\} \right) > h \right\}. \quad (4.9)$$

Then, $H(v, c_{P,h}(v))$ corresponds to the largest (for the inclusion order) half-space directed by v with P -mass at most h , which contains all the λv 's for λ large enough. We recall that for $c \in \mathbb{R}$, $H(v, c) = \{x \in \mathbb{R}^d \mid \langle x, v \rangle > c\}$.

Lemma 4.19. *Let $v \in S(0, 1)$ and $P \in \mathcal{P}^K(\mathbb{R}^d)$ for some $K > 0$. Assume that $P(\partial H(v, c_{P,h}(v))) = 0$. If $x_n = nv$ for all $n \in \mathbb{N}$, then for P -almost all $y \in \mathbb{R}^d$, we have:*

$$\lim_{n \rightarrow +\infty} \mathbb{1}_{\mathcal{B}(x_n, \delta_{P,h}(x_n))}(y) = \mathbb{1}_{H(v, c_{P,h}(v))}(y).$$

If $(x_n)_{n \in \mathbb{N}}$ is a sequence of \mathbb{R}^d such that $\lim_{n \rightarrow +\infty} d_{\overline{\mathbb{R}}^d}(x_n, v_\infty) = 0$, then, the result holds up to a subsequence.

The proof of Lemma 4.19 is given in the Appendix, Section 4.4.3.

For all $P \in \mathcal{P}_2(\mathbb{R}^d)$, we can generalize the definition of $\mathcal{P}_{x,h}(P)$, $P_{x,h}$, $m(P_{x,h})$, $v(P_{x,h})$ and $M(P_{x,h})$ to the elements $x = v_\infty \in \overline{\mathbb{R}}^d \setminus \mathbb{R}^d$ for all $v \in S(0, 1)$. Elements $P_{x,h}$ in $\mathcal{P}_{x,h}(P)$ are Borel probability measures $P_{x,h} = \frac{1}{h}Q$ with Q a sub-measure of P supported on $\overline{H}(v, c_{P,h}(v))$ and coinciding with P on $H(v, c_{P,h}(v))$.

Note that when P puts no mass on the hyperplane $\partial H(v, c_{P,h}(v))$, the set $\mathcal{P}_{v_\infty,h}(P)$ is reduced to the singleton $\{P_{v_\infty,h}\}$ with $P_{v_\infty,h}$ equal to $\frac{1}{h}P(B \cap H(v, c_{P,h}(v)))$ for all Borel sets B .

We define the set $\overline{\mathcal{M}}_h(P)$ by

$$\overline{\mathcal{M}}_h(P) = \left\{ m(P_{x,h}) \mid x \in \overline{\mathbb{R}}^d, P_{x,h} \in \mathcal{P}_{x,h}(P) \right\}. \quad (4.10)$$

Lemma 4.20. *Let $P \in \mathcal{P}^K(\mathbb{R}^d)$ for some $K > 0$. Then, the sets $\overline{\mathcal{M}}_h(P)$ and $\text{Conv}(\overline{\mathcal{M}}_h(P))$ are closed.*

The proof of Lemma 4.20 is deferred to the Appendix, in Section 4.4.3.

Intuitively, the distributions $P_{v_\infty, h}$ behave like extreme points of $\{P_{x, h} \mid x \in \mathbb{R}^d\}$. This intuition is formalized by the following Lemma.

Lemma 4.21. *Let $P \in \mathcal{P}^K(\mathbb{R}^d)$ for some $K > 0$. Then, the convex hull of $\overline{\mathcal{M}}_h(P)$ satisfies*

$$\text{Conv}(\overline{\mathcal{M}}_h(P)) = \bigcap_{v \in \mathbb{S}(0,1)} H^c(v, \langle m(P_{v_\infty, h}), v \rangle).$$

A straightforward consequence of Lemma 4.21 is the following.

Lemma 4.22. *Let $P \in \mathcal{P}^K(\mathbb{R}^d)$ for some $K > 0$ and $h, h' \in (0, 1)$ such that $h' \geq h$. Then,*

$$\text{Conv}(\overline{\mathcal{M}}_{h'}(P)) \subset \text{Conv}(\overline{\mathcal{M}}_h(P)).$$

The proof of Lemma 4.21 is to be found in Section 4.4.3. The proof of Lemma 4.22 is to be found in Section 4.4.3.

Since $\tilde{\mathcal{M}}_h(P)$ is compact, the proof of Lemma 4.21 can be adapted to the set $\tilde{\mathcal{M}}_h(P)$, and we get that:

$$\tilde{\mathcal{M}}_h(P) = \text{Conv}(\tilde{\mathcal{M}}_h(P)) = \bigcap_{v \in \mathbb{S}(0,1)} H^c(v, \langle m(P_{v_\infty, h}), v \rangle).$$

As a consequence, $h \mapsto \tilde{\mathcal{M}}_h(P)$ is non-increasing, see Lemma 4.22. We will see that according to Proposition 4.25 and Lemma 4.21, $\tilde{\mathcal{M}}_h(P)$ and $\overline{\mathcal{M}}_h(P)$ coincide when $P \in \mathcal{P}^K(\mathbb{R}^d)$ puts mass neither on balls nor on hyperplanes. Indeed, in this case the set $\overline{\mathcal{M}}_h(P)$ is convex, as will be proved in the following.

4.3.2 Convexity of the set of local means

The discrete case

We begin with the finite-sample case.

Points in a finite subset \mathbb{X} of \mathbb{R}^d are said to be *in general position* if any subset of \mathbb{X} with size at most $d + 1$ is a set of affinely independent points; see [BCY17][Section 3.1.4].

Lemma 4.23. *Let $P_n \in \mathcal{P}_n(\mathbb{R}^d)$. For $q \in \llbracket 1, n \rrbracket$ and $h = \frac{q}{n}$, set*

$$\hat{\mathcal{M}}_h(P_n) = \left\{ \bar{x} = \frac{1}{q} \sum_{p \in \text{NN}_{q, \mathbb{X}_n}(x)} p \mid x \in \overline{\mathbb{R}}^d, \text{NN}_{q, \mathbb{X}_n}(x) \in \mathcal{N}_{q, \mathbb{X}_n}(x) \right\},$$

with $\mathcal{N}_{q, \mathbb{X}_n}(x)$ the collection of all sets of q -nearest neighbors associated to x . Assume that the elements of $\hat{\mathcal{M}}_h(P_n)$ are in general position. Then, the set $\overline{\mathcal{M}}_{\frac{q}{n}}(P_n)$ is convex.

The proof of Lemma 4.23 is to be found in Section 4.4.3.

The general case with no mass on spheres nor hyperplanes

If $P \in \mathcal{P}^K(\mathbb{R}^d)$, for some $K > 0$, puts mass neither on boundary of balls, nor on hyperplanes, then the convexity of $\overline{\mathcal{M}}_h(P)$ may be deduced from the above Lemma 4.23 using the convergence of the empirical distribution P_n towards P in a probabilistic sense. This is summarized by the following Lemma.

Lemma 4.24. *Let $P \in \mathcal{P}^K(\mathbb{R}^d)$ be such that $P(S(x, r)) = 0$ and $P(\partial H(v, c)) = 0$ for all $x \in \mathbb{R}^d$, $r > 0$, $v \in S(0, 1)$ and $c \in \mathbb{R}$. Assume that with probability 1, for all $h = \frac{q}{n}$ for $n \in \mathbb{N}^*$ and $q \in \llbracket 1, n-1 \rrbracket$ large enough, the points in $\hat{\mathcal{M}}_h(P_n)$, defined in 4.23, are in general position.*

Set $h \in (0, 1)$ and $\theta \in \text{Conv}(\overline{\mathcal{M}}_h(P))$. Then, there are sequences $q_n \in \mathbb{N}$, $\alpha_n \rightarrow 0$, $P_n \in \mathcal{P}_n(\mathbb{R}^d)$ and $y_n \in \text{Conv}(\overline{\mathcal{M}}_{\frac{q_n}{n}}(P_n))$ such that

- i) $\frac{q_n}{n} \rightarrow h$,*
- ii) $\|y_n - \theta\| \leq \alpha_n$,*
- iii) $\sup_{x \in \overline{\mathbb{R}}^d} \|m(P_n, x, \frac{q_n}{n}) - m(P_{x,h})\| \leq \alpha_n$.*

Lemma 4.24 follows from probabilistic arguments when \mathbb{X}_n is sampled at random. Its proof can be found in Section 4.4.3. Equipped with Lemma 4.24, we can prove the convexity of $\overline{\mathcal{M}}_h(P)$.

Proposition 4.25. *Let $P \in \mathcal{P}^K(\mathbb{R}^d)$ be such that $P(S(x, r)) = 0$ and $P(\partial H(v, c)) = 0$ for all $x \in \mathbb{R}^d$, $r > 0$, $v \in S(0, 1)$ and $c \in \mathbb{R}$. Assume that with probability 1, for all $h = \frac{q}{n}$ for $n \in \mathbb{N}^*$ and $q \in \llbracket 1, n-1 \rrbracket$ large enough, the points in $\hat{\mathcal{M}}_h(P_n)$, defined in 4.23, are in general position.*

Then, for all $h \in (0, 1)$, $\overline{\mathcal{M}}_h(P)$ is convex.

Proof of Proposition 4.25. Let $\theta \in \text{Conv}(\overline{\mathcal{M}}_h(P))$, P_n, q_n, α_n, y_n be as in Lemma 4.24, and for short let $h_n = \frac{q_n}{n}$.

Since $\overline{\mathcal{M}}_{h_n}(P_n)$ is convex according to Lemma 4.23, there is a sequence $(x_n)_{n \geq N}$ in $\overline{\mathbb{R}}^d$ such that $y_n = m(P_n, x_n, h_n)$. If $(x_n)_{n \geq N}$ is bounded, then up to a subsequence we have $x_n \rightarrow x$, for some $x \in \mathbb{R}^d$. Otherwise, considering $\frac{x_n}{\|x_n\|}$, up to a subsequence we get $x_n \rightarrow v_\infty$. In any case $x_n \rightarrow x$, for $x \in \overline{\mathbb{R}}^d$. Combining Lemma 4.24, Lemma 4.19 and the dominated convergence lemma yields $\theta = m(P_{x,h})$. Thus, $\theta \in \overline{\mathcal{M}}_h(P)$. \square

We have proved convexity of the ‘‘closure’’ of the set of local means when P puts mass neither on spheres nor on hyperplanes. With some more work, we hope that we may not need the assumption about the local means in general position.

4.3.3 Duality, set of local means and distance-to-measure

For $P \in \mathcal{P}(\mathbb{R}^d)$ and $h \in [0, 1]$, set $\psi_{P,h}$ to be the function defined on \mathbb{R}^d by

$$\psi_{P,h} : x \mapsto \frac{1}{2} (\|x\|^2 - d_{P,h}^2(x)). \tag{4.11}$$

In [CCSM11, Proposition 3.6], Chazal et al. prove that $\psi_{P,h}$ is a convex function and that its sub-gradients are local means. Moreover, when P puts mass neither on $\partial \mathcal{B}(x, \delta_{P,h}(x))$ nor on $\partial H(v, c_{P,h}(v))$ for $x \in \mathbb{R}^d$ and $v \in S(0, 1)$, $\psi_{P,h}$ is differentiable with gradient at x given by $m(P_{x,h})$. With the additional assumption of strict convexity for $\psi_{P,h}$, we show that Legendre duality applied to the convex function of Legendre type $\psi_{P,h}$ provides another proof of convexity for the set of local means.

Next, we relate the functions $\psi_{P,h}$ and $\omega_{P,h}$ defined in Section 4.1 in (4.3). More precisely, we prove that $\psi_{P,h}^*$ can be expressed for all $\theta \in \mathbb{R}^d$ as:

$$\psi_{P,h}^*(\theta) = \frac{1}{2} (\omega_{P,h}^2(\theta) + \|\theta\|^2).$$

We also note that when P puts mass neither on $\partial\mathcal{B}(x, \delta_{P,h}(x))$ nor on $\partial\mathbb{H}(v, c_{P,h}(v))$ for $x \in \mathbb{R}^d$ and $v \in \mathbb{S}(0, 1)$, the Bregman divergence associated to the convex and differentiable function $\psi_{P,h}$ satisfies for all $x, y \in \mathbb{R}^d$:

$$d_{\psi_{P,h}}(x, y) = \frac{1}{2} (\|x - m(P_{y,h})\|^2 + v(P_{y,h}) - d_{P,h}^2(x)).$$

As a consequence, it turns out that the k -PDTM defined in Section 4.1 is built from a codebook minimizing the Bregman k -divergence-to-measure defined in Section 3.1, associated to the Bregman divergence $d_{\psi_{P,h}}$, that is to say, from an optimal codebook for the Bregman quantization cost function with the Bregman divergence associated to $\psi_{P,h}$.

The set of local means, a set of sub-gradients of a convex function

Given $f : \mathbb{R}^d \rightarrow [-\infty, +\infty]$ a convex function, a *subgradient* of f at a point $x \in \mathbb{R}^d$ is an element $c \in \mathbb{R}^d$ such that

$$f(z) \geq f(x) + \langle c, z - x \rangle, \forall z \in \mathbb{R}^d. \quad (4.12)$$

The set of all subgradients of a convex function f at a point x is called the *subgradient set* and denoted by $\partial_x f$.

It is well-known that when f is strictly convex, equality holds in (4.12) if and only if $x = z$. We need to recall the following well-known theorem.

Theorem 4.26 ([Roc70, Theorem 25.1]). *Let $f : \mathbb{R}^d \rightarrow [-\infty, +\infty]$ be a convex function, and let x be a point where f is finite. If f is differentiable at x , then $\nabla f(x)$ is the unique subgradient of f at x . Conversely, if f has a unique subgradient at x , then f is differentiable at x .*

Chazal et al. noticed in [CCSM11] that the distance-to-measure P is semiconcave, and its set of local means plays the role of a subgradient set in the following sense.

Proposition 4.27 ([CCSM11, Proposition 3.3]). *The map*

$$\psi_{P,h} : x \mapsto \frac{1}{2} (\|x\|^2 - d_{P,h}^2(x)) \quad (4.13)$$

is convex. Moreover, its subgradient set is given at any point $x \in \mathbb{R}^d$ by:

$$\partial_x \psi = \{m(P_{x,h}) \mid P_{x,h} \in \mathcal{P}_{x,h}(P)\}.$$

A direct consequence of Proposition 4.27 has been very useful in this chapter.

Proposition 4.28 ([CCSM11, Proposition 3.6]). *For all $x, y \in \mathbb{R}^d$ and $P_{x,h} \in \mathcal{P}_{x,h}(P)$,*

$$d_{P,h}^2(y) - \|y\|^2 \leq d_{P,h}^2(x) - \|x\|^2 - 2\langle y - x, m(P_{x,h}) \rangle,$$

with equality if and only if $P_{x,h} \in \mathcal{P}_{y,h}(P)$.

Another consequence of Proposition 4.27 is the fact that the DTM is 1-Lipschitz.

Note that under the assumption $P(\partial\mathcal{B}(x, \delta_{P,h}(x))) = 0$ for all $x \in \mathbb{R}^d$, all $\mathcal{P}_{x,h}(P)$ are reduced to singletons. Thus, according to Theorem 4.26, the map $\psi_{P,h}$ is differentiable, with gradient $\nabla\psi_{P,h}(x) = m(P_{x,h})$ where $P_{x,h}$ is the restriction of the measure P to the ball $\mathcal{B}(x, \delta_{P,h}(x))$.

A proof of the convexity of the set of local means based on Legendre duality

First, we recall the notions of Legendre duality and conjugate functions as defined in [Roc70] or [BMDG05].

The *effective domain* of a convex function $f : \mathbb{R}^d \rightarrow \llbracket -\infty, +\infty \rrbracket$, denoted by $\text{dom}(f)$, is the set of points at which f is not equal to $+\infty$. A convex function is *proper* if $\text{dom}(\psi)$ is non-empty and:

$$\forall x \in \text{dom}(\psi), \psi(x) > -\infty.$$

A convex function is *closed* if it is lower semi-continuous, that is:

$$\forall x_0 \in \text{dom}(\psi), \liminf_{x \rightarrow x_0} \psi(x) \geq \psi(x_0).$$

Definition 4.29 ([Roc70]). Let ψ be a proper, closed convex function with $\Theta = \text{dom}(\psi)^\circ$. The pair (Θ, ψ) is called a *convex function of Legendre type* or a Legendre function if the following properties are satisfied:

- (i) Θ is non-empty
- (ii) ψ is strictly convex and differentiable on Θ
- (iii) $\forall \theta_b \in \partial\Theta, \lim_{\theta \rightarrow \theta_b, \theta \in \Theta} \|\nabla\psi(\theta)\| = \infty$.

In particular, whenever ψ is a real-valued convex function defined on \mathbb{R}^d , (\mathbb{R}^d, ψ) is a convex function of Legendre type if and only if ψ is strictly convex and differentiable on \mathbb{R}^d . Indeed, the remaining assumptions are automatically satisfied, for instance closedness follows from the differentiability of ψ on \mathbb{R}^d ; see Theorem 4.26.

We now define the notion of conjugate function.

Definition 4.30 (Theorem 12.2 [Roc70]). Let ψ be a real-valued function on \mathbb{R}^d . Then its *conjugate function* ψ^* is given by

$$\psi^*(t) = \sup_{\theta \in \text{dom}(\psi)} \{\langle t, \theta \rangle - \psi(\theta)\}. \quad (4.14)$$

If ψ is a proper closed convex function, ψ^* is also a proper closed convex function and $\psi_{P,h}^{**} = \psi_{P,h}$.

Note that since the DTM is 1-Lipschitz, $\psi_{P,h}$ is continuous thus closed. In addition, $\psi_{P,h}$ is proper. So, according to Definition 4.30, its conjugate function $\psi_{P,h}^*$ is also a proper closed convex function, and moreover, $\psi^{**} = \psi$.

The following theorem will be used to derive another proof for the convexity of $\mathcal{M}_h(P)$ under the assumption of strict convexity for measures that put mass neither on spheres nor on hyperplanes.

Theorem 4.31 (Theorem 26.5 [Roc70]). Let ψ be a real-valued proper closed convex function with conjugate function ψ^* . Let $\Theta = \text{dom}(\psi)^\circ$ and $\Theta^* = \text{dom}(\psi^*)^\circ$. If (Θ, ψ) is a convex function of Legendre type, then

- (i) (Θ^*, ψ^*) is a convex function of Legendre type,
- (ii) (Θ, ψ) and (Θ^*, ψ^*) are Legendre dual to each other,
- (iii) The gradient function $\nabla\psi : \Theta \rightarrow \Theta^*$ is a one-to-one function from the open convex set Θ onto the open convex set Θ^* ,
- (iv) The gradient functions $\nabla\psi, \nabla\psi^*$ are continuous, and $\nabla\psi^* = (\nabla\psi)^{-1}$.

The convexity of $\mathcal{M}_h(P)$ comes from the fact that the set of gradients of $\psi_{P,h}$ is exactly equal to the set of local means.

Legendre duality provides a new set of assumptions under which the set of local means is convex.

Proposition 4.32. *Let $h \in (0, 1)$ and $P \in \mathcal{P}(\mathbb{R}^d)$. If P satisfies the following assumptions:*

(i) $\forall x \in \mathbb{R}^d, P(\partial \mathcal{B}(x, \delta_{P,h}(x))) = 0,$

(ii) $\psi_{P,h}$ is strictly convex,

then $\mathcal{M}_h(P)$ is convex.

Proof. The function $\psi_{P,h}$ defined by (4.11) is a real-valued function defined on \mathbb{R}^d which is convex, according to Proposition 4.27. Thus, $(\mathbb{R}^d, \psi_{P,h})$ is of Legendre type if and only if $\psi_{P,h}$ is differentiable and strictly convex on \mathbb{R}^d . The differentiability comes when P puts no mass on balls of P -mass h . Moreover, Proposition 4.27 and definition of $\mathcal{P}_{x,h}(P)$ yield that in this case $\mathcal{M}_h(P)$ is exactly equal to the set of gradients of $\psi_{P,h}$. The convexity of $\mathcal{M}_h(P)$ under the additional assumption of strict convexity for $\psi_{P,h}$ follows from Theorem 4.31. \square

Proposition 4.32 is in a sense more general than Proposition 4.25 since P does not need to be supported on a compact set, but more restrictive since it requires strict convexity for $\psi_{P,h}$ on \mathbb{R}^d . The assumption of strict convexity is obviously not satisfied when a ball $\mathcal{B}(x, r)$ is of P -mass h and is separated from the remaining support of P in the sense that for some $\epsilon > 0$, $P(\mathcal{B}(x, r + \epsilon) \setminus \mathcal{B}(x, r)) = 0$. Indeed, all points y in $\mathcal{B}(x, \epsilon)$ will share the same gradient for $\psi_{P,h}$: $\nabla \psi_{P,h}(y) = m(P_{x,h})$.

A connection between the k -PDTM (Section 4.1) and the clustering with Bregman divergences (Section 3.1)

The function $\omega_{P,h}$ defined by Equation (4.3) is finite on the set $\text{Conv}(\overline{\mathcal{M}}_h(P))$ and infinite outside.

Lemma 4.33. *Let $P \in \mathcal{P}^K(\mathbb{R}^d)$ for some $K > 0$. Then $\theta \in \text{Conv}(\overline{\mathcal{M}}_h(P))$ if and only if $\omega_{P,h}(\theta) < +\infty$.*

The proof of this lemma is deferred to the Appendix in Section 4.4.3.

This property is related to the fact that the conjugate function of $\psi_{P,h}, \psi_{P,h}^*$ can be rewritten as a function of $\omega_{P,h}$ as follows.

Proposition 4.34. *For all measures $P \in \mathcal{P}(\mathbb{R}^d)$,*

$$\psi_{P,h}^*(\theta) = \frac{1}{2} (\omega_{P,h}^2(\theta) + \|\theta\|^2), \forall \theta \in \mathbb{R}^d.$$

This function is a proper closed convex function and satisfies at all $x \in \mathbb{R}^d$,

$$\psi_{P,h}^{**}(x) = \frac{1}{2} (\|x\|^2 - d_{P,h}^2(x)).$$

Proof. The conjugate function of $\psi_{P,h}$ can be expressed as, for $\theta \in \mathbb{R}^d$,

$$\begin{aligned} \psi_{P,h}^*(\theta) &= \sup_{x \in \mathbb{R}^d} \langle x, \theta \rangle - \psi_{P,h}(x) \\ &= \frac{1}{2} \sup_{x \in \mathbb{R}^d} \{ \|\theta\|^2 - \|\theta - x\|^2 + d_{P,h}^2(x) \} \\ &= \frac{1}{2} (\omega_{P,h}^2(\theta) + \|\theta\|^2). \end{aligned}$$

The remaining statements are direct consequences of the fact that the DTM is Lipschitz and Definition 4.30 as aforementioned. \square

We can make another connection between the convex function $\psi_{P,h}$ and the k -PDTM defined in Section 4.1, but also with the k -Bregman divergence-to-measure, defined in Section 3.1 as follows. The Bregman divergence associated to $\psi_{P,h}$ satisfies for all $x, y \in \mathbb{R}^d$:

$$\begin{aligned} d_{\psi_{P,h}}(x, y) &= \psi_{P,h}(x) - \psi_{P,h}(y) - \langle \nabla \psi_{P,h}(y), x - y \rangle \\ &= \frac{1}{2} (\|x\|^2 - d_{P,h}^2(x) - \|y\|^2 + d_{P,h}^2(y) - 2\langle m(P_{y,h}), x - y \rangle) \\ &= \frac{1}{2} (\|x\|^2 - d_{P,h}^2(x) + v(P_{y,h}) + \|m(P_{y,h})\|^2 - 2\langle m(P_{y,h}), x \rangle) \\ &= \frac{1}{2} (\|x - m(P_{y,h})\|^2 + v(P_{y,h}) - d_{P,h}^2(x)). \end{aligned}$$

In particular, the sets of points c_1, c_2, \dots, c_k on \mathbb{R}^d that minimize the function

$$(c_1, c_2, \dots, c_k) \mapsto P \min_{i \in [1,k]} d_{\psi_{P,h}}(u, c_i) \tag{4.15}$$

are exactly the set of points that minimize the function

$$(c_1, c_2, \dots, c_k) \mapsto P \min_{i \in [1,k]} \|u - m(P_{c_i,h})\|^2 + v(P_{c_i,h}). \tag{4.16}$$

Thus, when P does not charge boundary of balls of P -mass h , any k -PDTM obtained from a minimizer of (4.16) (see Section 4.1) corresponds to a clustering in the sense of the minimum of the Bregman k -divergence to the measure P with parameter $h = 1$ as studied in Section 3.1, see also [BMDG05].

In Section 4.2, we have adapted the concept of trimmed k -PDTM by applying the methods in Section 3.1 to the Bregman divergence associated with $\psi_{P,h}$. Such a concept furnishes a method of quantization even less sensible to outliers.

4.4 Proofs

4.4.1 Proofs for Section 4.1

Proof of Lemma 4.1

Since $\mathcal{P}_h(\mathbb{R}^d)$ is convex, it is immediate that $\tilde{\mathcal{M}}_h(P)$ is convex. Let $(m(Q_n))_{n \in \mathbb{N}}$ be a sequence in $\tilde{\mathcal{M}}_h(P)$ that converges towards t , with $Q_n \in \mathcal{P}_h(\mathbb{R}^d)$. The set $\mathcal{P}_h(\mathbb{R}^d)$ is tight. Thus, according to Prokhorov's Theorem, up to a subsequence, there exists Q a Borel positive measure on \mathbb{R}^d such that Q_n converges weakly to Q . Note that Q belongs to $\mathcal{P}_h(\mathbb{R}^d)$. Indeed, the dominated convergence Lemma applied to the continuous constant function equal to 1 implies that $Q(\mathbb{R}^d) = h$. Also, the Portmanteau Lemma entails that $Q(O) \leq P(O)$ for all open set O . Also, since Euclidean \mathbb{R}^d is a Polish space, [AB06, p.438] states that the measures Q and P are regular measures. Thus, Q is a sub-measure of P , and $Q \in \mathcal{P}_h(\mathbb{R}^d)$.

Let $\varepsilon > 0$, and for $M > 0$, $u \in \mathbb{R}^d$, denote by $u \wedge M$ the vector $(u_1 \wedge M, \dots, u_d \wedge M)$, where $u_i \wedge M = u_i$ when $u_i \in [-M, M]$ and M otherwise. Then there exists $M_\varepsilon > 0$ such that for every n

$$\begin{aligned} \|Q_n(u \wedge M_\varepsilon) - Q_n u\| &\leq P(\|u\| \mathbb{1}_{\|u\|_\infty > M_\varepsilon}) \\ &\leq \frac{\varepsilon}{2}. \end{aligned}$$

The same inequality holds for Q . On the other hand, since $u \mapsto u \wedge M_\varepsilon$ is bounded and continuous, we have, for n large enough,

$$\|Q_n(u \wedge M_\varepsilon) - Q(u \wedge M_\varepsilon)\| \leq \frac{\varepsilon}{2}.$$

Thus, taking into account that $Q_n u/h \rightarrow t$, we have, for every $\varepsilon > 0$, $\|t - \frac{1}{h}Q u\| \leq \frac{3\varepsilon}{h}$. Hence $t \in \tilde{\mathcal{M}}_h(P)$.

Proof of Lemma 4.4

According to Proposition 4.18, we may write

$$\omega_{P,h}^2(\theta) = \sup_{x \in \mathbb{R}^d} \inf_{Q \in \mathcal{P}_h(P)} g(x, Q),$$

where $g(x, Q) = M(Q) - \|\theta\|^2 + 2 \langle x, \theta - m(Q) \rangle$. First note that, according to Prokhorov's theorem, $\mathcal{P}_h(P)$ is a compact vector space (equipped with the weak convergence metric). As proved in Lemma 4.1, $Q \mapsto m(Q)$ is continuous on $\mathcal{P}_h(P)$. Now let $Q_n \rightarrow Q$ in distribution. Then, for every $M > 0$, $Q_n \|u\|^2 \wedge M \rightarrow Q \|u\|^2 \wedge M$. Since $\sup_n |Q_n(\|u\|^2 \wedge M - \|u\|^2)| \leq (P\|u\|^2 \mathbb{1}_{\|u\| > M})/h$, we deduce as well that $M(Q_n) \rightarrow M(Q)$. Thus, for every $x \in \mathbb{R}^d$, $g(x, \cdot)$ is continuous, and linear.

On the other hand, for every Q in $\tilde{\mathcal{M}}_h(P)$, $g(\cdot, Q)$ is linear and continuous. Thus, according to Sion's theorem [Kom88], we may write

$$\omega_{P,h}^2(\theta) = \inf_{Q \in \mathcal{P}_h(P)} \sup_{x \in \mathbb{R}^d} M(Q) - \|\theta\|^2 + 2 \langle x, \theta - m(Q) \rangle. \quad (4.17)$$

Thus $\omega_{P,h}^2(\theta) < \infty$ is equivalent to $\theta \in \tilde{\mathcal{M}}_h(P)$. Now let θ be in $\tilde{\mathcal{M}}_h(P)$. According to (4.17), we may write

$$\omega_{P,h}^2(\theta) = \inf_{Q \in \mathcal{P}_h(P), m(Q)=\theta} M(Q) - \|\theta\|^2 = \inf_{Q \in \mathcal{P}_h(P), m(Q)=\theta} v(Q).$$

Since $Q \mapsto v(Q)$ is continuous on $\mathcal{P}_h(P)$ and $\mathcal{P}_h(P) \cap m^{-1}(\{\theta\})$ is compact, there exists Q such that $m(Q) = \theta$ and $\omega_{P,h}^2(\theta) = v(Q)$.

Let $f : \mathbb{R}^{d(k)} \rightarrow \mathbb{R}$ such that $f(\tau_1, \dots, \tau_k) = P \min_{i \in [1,k]} \|u - \tau_i\|^2 + \omega_{P,h}^2(\tau_i)$. Since $\tau \mapsto \omega_{P,h}^2(\tau) = \sup_x d_{P,h}^2(x) - \|x - \tau\|^2$ is lower semi-continuous, Fatou's Lemma ensures that f is lower semi-continuous. Thus $f|_{\tilde{\mathcal{M}}_h(P)^{(k)}}$ has a minimum, which is a minimum of f according to the first part.

Proof of Lemma 4.8

For short, we use the notation $m_i = m(\tilde{P}_{i,h})$, $v_i = v(\tilde{P}_{i,h})$ and $Q(du) f(u)$ for the expectation of f with respect to the Borel probability measure $\frac{Q}{Q(\mathbb{R}^d)} \in \mathcal{P}(\mathbb{R}^d)$. Then, a bias-variance decomposition yields

$$\begin{aligned} R(P_1, \dots, P_k) &= P(du) \min_{i \in [1,k]} P_i(dz) \|u - z\|^2 \\ &= \sum_{i=1}^k \tilde{P}_{i,h}(du) P_i(dz) \|u - z\|^2 \\ &= \sum_{i=1}^k P_i(dz) \tilde{P}_{i,h}(\mathbb{R}^d) [\|z - m_i\|^2 + v_i] \\ &\geq \sum_{i=1}^k \tilde{P}_{i,h}(\mathbb{R}^d) P_{m_i,h}(dz) (\|z - m_i\|^2 + v_i), \end{aligned}$$

where $P_{m_i,h} \in \mathcal{P}_{m_i,h}(P)$ and equality holds if and only if $P_i \in \mathcal{P}_{m_i,h}(P)$, according to Proposition 4.1. Thus, we may write

$$\begin{aligned} R(P_1, \dots, P_k) &\geq \sum_{i=1}^k \tilde{P}_{m_i,h}(du) P_{m_i,h}(dz) [\|z - u\|^2] \\ &= \sum_{i=1}^k \tilde{P}_{m_i,h}(du) [\|m(P_{m_i,h}) - u\|^2 + v(P_{m_i,h})] \\ &= R(P_{m_1,h}, \dots, P_{m_k,h}). \end{aligned}$$

Proof of Corollary 4.11

The proof of Corollary 4.11 is based on the following bounds, in the case where P is absolutely continuous with respect to the Lebesgue measure, with density f satisfying $0 < f_{min} \leq f \leq f_{max}$.

$$f_M^{-1}(k) \leq 2K\sqrt{d}k^{-1/d} \tag{4.18}$$

$$\zeta_{P,h}(f_M^{-1}(k)) \leq C_{f_{max},K,d,h}k^{-1/d}. \tag{4.19}$$

The first equation proceeds from the following. Since $M \subset \mathcal{B}(0, K)$, for any $\varepsilon > 0$ we have

$$f_M(\varepsilon) \leq f_{\mathcal{B}(0,K)}(\varepsilon) \leq \left(\frac{2K\sqrt{d}}{\varepsilon}\right)^d.$$

Hence (4.18). To prove the second inequality, we will use the following Lemma.

Lemma 4.35. *Suppose that P has a density f satisfying $0 < f_{min} \leq f \leq f_{max}$. Let x, y be in M , and denote by $\delta = \|x - y\|$. Then*

$$\|m(P_{x,h}) - m(P_{y,h})\| \leq \frac{2dK^{d+1}\omega_d}{h} \left(1 + \delta \left(\frac{f_{max}\omega_d}{h}\right)^{1/d}\right)^{d-1} \left(\frac{f_{max}\omega_d}{h}\right)^{1/d} \delta,$$

with ω_d the Lebesgue volume of the ball $\mathcal{B}(0, 1)$ in \mathbb{R}^d .

Proof of Lemma 4.35. Since P has a density, $P\partial\mathcal{B}(x, \delta_{P,h}(x)) = P\partial\mathcal{B}(y, \delta_{P,h}(y)) = 0$. We deduce that $P_{x,h} = \frac{1}{h}P|_{\mathcal{B}(x, \delta_{P,h}(x))}$ and $P_{y,h} = \frac{1}{h}P|_{\mathcal{B}(y, \delta_{P,h}(y))}$. Without loss of generality, assume that $\delta_{P,h}(x) \geq \delta_{P,h}(y)$. Then $\mathcal{B}(y, \delta_{P,h}(y)) \subset \mathcal{B}(x, \delta_{P,h}(x) + \delta)$. We may write

$$\begin{aligned} \|m(P_{x,h}) - m(P_{y,h})\| &= \frac{1}{h} \left\| Pu \left(\mathbb{1}_{\mathcal{B}(x, \delta_{P,h}(x))}(u) - \mathbb{1}_{\mathcal{B}(y, \delta_{P,h}(y))}(u) \right) \right\| \\ &\leq \frac{K}{h} P \left| \mathbb{1}_{\mathcal{B}(x, \delta_{P,h}(x))}(u) - \mathbb{1}_{\mathcal{B}(y, \delta_{P,h}(y))}(u) \right| \\ &= 2\frac{K}{h} P(\mathcal{B}(y, \delta_{P,h}(y)) \setminus (\mathcal{B}(x, \delta_{P,h}(x)) \cap \mathcal{B}(y, \delta_{P,h}(y)))) \\ &\leq 2\frac{K}{h} P(\mathcal{B}(x, \delta_{P,h}(x) + \delta) \cap \mathcal{B}(x, \delta_{P,h}(x))^c) \\ &= 2\frac{K}{h} \omega_d \left[(\delta_{P,h}(x) + \delta)^d - \delta_{P,h}(x)^d \right] \\ &\leq 2\frac{K^{d+1}\omega_d}{h} \left[\left(1 + \frac{\delta}{\delta_{P,h}(x)}\right)^d - 1 \right]. \end{aligned}$$

Since $(1 + v)^d \leq 1 + d(1 + v)^{d-1}v$, for $v \geq 0$, and $\delta_{P,h}(x) \geq \left(\frac{h}{f_{max}\omega_d}\right)^{1/d}$, the result follows. □

Hence (4.19). The result of Corollary 4.11 follows.

Proof of Corollary 4.12

Without loss of generality we assume that N is connected. Since P has a density with respect to the volume measure on N , we have $P(N^\circ) = 1$. Thus we take $M = N^\circ$, that is the set of interior points. Since P satisfies a (cf_{min}, d') -standard assumption, we have

$$f_M(\varepsilon) \leq \frac{2^{d'}}{cf_{min}} r^{-d'},$$

according to [CGLM15, Lemma 10]. Hence $f_M^{-1}(k) \leq C_{f_{min}, N} k^{-1/d'}$. It remains to bound the continuity modulus of $x \mapsto m(P_{x,h})$. For any x in M , since $P(\partial N) = 0$ and P has a density with respect to the volume measure on N , we have $P_{x,h} = P|_{\mathcal{B}(x,h)}$. Besides, since for all $r > 0$ $P(\mathcal{B}(x,r)) \geq cf_{min}r^{d'}$, we may write $\delta_{P,h}(x) \leq c_{N,f_{min}} h^{1/d'} \leq \rho/12$, for h small enough. Now let x and y be in M so that $\|x - y\| = \delta \leq \rho/12$, and without loss of generality assume that $\delta_{P,h}(x) \geq \delta_{P,h}(y)$. Then, proceeding as in the proof of Lemma 4.35, it comes

$$\|m(P_{x,h}) - m(P_{y,h})\| \leq \frac{2K}{h} P(\mathcal{B}(x, \delta_{P,h}(x) + \delta) \cap \mathcal{B}(x, \delta_{P,h}(x))^c).$$

Since $\delta_{P,h}(x) + \delta \leq \rho/6$, for any u in $\mathcal{B}(x, \delta_{P,h}(x) + \delta) \cap M$ we may write $u = \exp_x(rv)$, where $v \in T_x M$ with $\|v\| = 1$ and $r = d_N(u, x)$ is the geodesic distance between u and x (see [Fed59, Theorem 4.18] or e.g., [EC17, Proposition 25]). Note that, according to [EC17, Proposition 26], for any u_1 and u_2 such that $\|u_1 - u_2\| \leq \rho/4$,

$$\|u_1 - u_2\| \leq d_N(u_1, u_2) \leq 2\|u_1 - u_2\|. \tag{4.20}$$

Now let p_1, \dots, p_m be a δ -covering set of the sphere $S(x, \delta_{P,h}(x)) = \{u \in M \mid \|x - u\| = \delta_{P,h}(x)\}$. According to (4.20), we may choose $m \leq c_{d'} \delta_{P,h}(x)^{d'-1} \delta^{-(d'-1)}$.

Now, for any u such that $u \in M$ and $\delta_{P,h}(x) \leq \|x - u\| \leq \delta_{P,h}(x) + \delta$, there exists $t \in \mathcal{S}_{x, \delta_{P,h}(x)}$ such that $\|t - u\| \leq 2\delta$. Hence

$$P(\mathcal{B}(x, \delta_{P,h}(x) + \delta) \cap \mathcal{B}(x, \delta_{P,h}(x))^c) \leq \sum_{j=1}^m P(\mathcal{B}(p_j, 2\delta)).$$

Now, for any j , since $2\delta \leq \rho/6$, in local polar coordinates around p_j we may write, using (4.20) again,

$$\begin{aligned} P(\mathcal{B}(p_j, 2\delta)) &\leq \int_{r,v \mid \exp_{p_j}(rv) \in M, r \leq 4\delta} f(r, v) J(r, v) dr dv \\ &\leq f_{max} \int_{r,v \mid r \leq 4\delta} J(r, v) dr dv \end{aligned}$$

where $J(r, v)$ denotes the Jacobian of the volume form. According to [EC17, Proposition 27], we have $J(r, v) \leq C_{d'} r^{d'}$. Hence $P(\mathcal{B}(p_j, 2\delta)) \leq C_{d'} f_{max} \delta^{d'}$. We may conclude

$$\begin{aligned} \|m(P_{x,h}) - m(P_{y,h})\| &\leq \frac{2K}{h} m C_{d'} f_{max} \delta^{d'} \\ &\leq C_{N, f_{max}, f_{min}} \delta. \end{aligned}$$

Choosing k large enough so that $f_M^{-1}(k) \leq C_{f_{min}, N} k^{-1/d'} \leq \rho/12$ gives the result of Corollary 4.12.

Proof of Proposition 4.13

For all $x \in \text{Supp}(P)$,

$$\begin{aligned} d_{Q,h,k}^2(x) - d_{P,h}^2(x) &= d_{Q,h,k}^2(x) - d_{Q,h}^2(x) + d_{Q,h}^2(x) - d_{P,h}^2(x) \\ &\geq -\|d_{P,h}^2 - d_{Q,h}^2\|_{\infty, \text{Supp}(P)}. \end{aligned}$$

Thus, $(d_{Q,h,k}^2 - d_{P,h}^2)_- \leq \|d_{P,h}^2 - d_{Q,h}^2\|_{\infty, \text{Supp}(P)}$ on $\text{Supp}(P)$, where $f_- : x \mapsto f(x)\mathbb{1}_{f(x) \leq 0}$ denotes the negative part of any function f on \mathbb{R}^d . Then,

$$\begin{aligned} P|d_{Q,h,k}^2 - d_{P,h}^2|(u) &= P d_{Q,h,k}^2(u) - d_{P,h}^2(u) + 2(d_{Q,h,k}^2(u) - d_{P,h}^2(u))_- \\ &\leq P\Delta(u) + P d_{P,h,k}^2(u) - d_{P,h}^2(u) + 2\|d_{P,h}^2 - d_{Q,h}^2\|_{\infty, \text{Supp}(P)}. \end{aligned}$$

with $\Delta = d_{Q,h,k}^2 - d_{P,h,k}^2$. We can bound $P\Delta(u)$ from above. Let $s \in \mathcal{O}pt(P, h, k) \cap \overline{\mathcal{B}}(0, K)^{(k)}$ such that for all $i \in \llbracket 1, k \rrbracket$ such that $\tilde{P}_{s_i, h}(\mathbb{R}^d) \neq 0$, $s_i = m(\tilde{P}_{s_i, h})$. Such an s exists according to Lemma 4.8 and Lemma 4.4. Set $f_{Q,t}(x) = -2\langle x, m(Q_{t,h}) \rangle + M(Q_{t,h})$ for $t \in \mathbb{R}^d$, and let $t \in \mathcal{O}pt(Q, h, k)$.

$$\begin{aligned} P\Delta(u) &= P \min_{i \in \llbracket 1, k \rrbracket} f_{Q,t_i}(u) - \min_{i \in \llbracket 1, k \rrbracket} f_{P,s_i}(u) \\ &\leq (P - Q) \min_{i \in \llbracket 1, k \rrbracket} f_{Q,t_i}(u) + (Q - P) \min_{i \in \llbracket 1, k \rrbracket} f_{Q,s_i}(u) + P \min_{i \in \llbracket 1, k \rrbracket} f_{Q,s_i}(u) - \min_{i \in \llbracket 1, k \rrbracket} f_{P,s_i}(u). \end{aligned}$$

For any transport plan π between P and Q ,

$$\begin{aligned} (P - Q) \min_{i \in \llbracket 1, k \rrbracket} f_{Q,t_i}(u) &= \mathbb{E}_{(X,Y) \sim \pi} \left[\min_{i \in \llbracket 1, k \rrbracket} -2\langle X, m(Q_{t_i,h}) \rangle + M(Q_{t_i,h}) - \min_{i \in \llbracket 1, k \rrbracket} -2\langle Y, m(Q_{t_i,h}) \rangle + M(Q_{t_i,h}) \right] \\ &\leq 2\mathbb{E}_{(X,Y) \sim \pi} \left[\sup_{t \in \mathbb{R}^d} \langle Y - X, m(Q_{t,h}) \rangle \right]. \end{aligned}$$

Thus, $(P - Q) \min_{i \in \llbracket 1, k \rrbracket} f_{Q,t_i}(u) \leq 2W_1(P, Q) \sup_{t \in \mathbb{R}^d} \|m(Q_{t,h})\|$, after taking for π the optimal transport plan for the L_1 -Wasserstein distance (noted W_1) between P and Q .

Also note that $P \min_{i \in \llbracket 1, k \rrbracket} f_{Q,s_i}(u) - \min_{i \in \llbracket 1, k \rrbracket} f_{P,s_i}(u)$ is bounded from above by

$$\begin{aligned} &\leq \sum_{i=1}^k \tilde{P}_{s_i, h}(-2\langle u, m(Q_{s_i,h}) \rangle + M(Q_{s_i,h})) - (-2\langle u, m(P_{s_i,h}) \rangle + M(P_{s_i,h})) \\ &= \sum_{i=1}^k \tilde{P}_{s_i, h} 2\langle u - s_i, m(P_{s_i,h}) - m(Q_{s_i,h}) \rangle + d_{Q,h}^2(s_i) - d_{P,h}^2(s_i) \\ &\leq \|d_{P,h}^2 - d_{Q,h}^2\|_{\infty, \mathcal{B}(0,K)} + 2 \sum_{i=1}^k \tilde{P}_{s_i, h}(\mathbb{R}^d) \langle m(\tilde{P}_{s_i, h}) - s_i, m(P_{s_i,h}) - m(Q_{s_i,h}) \rangle. \end{aligned}$$

Since $s_i = m(\tilde{P}_{s_i, h})$, the result follows.

Proof of Proposition 4.14

Let $\Delta_{\infty, K}$ denote $\sup_{x \in M} d_{Q,h,k}(x)$, and let $x \in M$ achieving the maximum distance. Since $d_{Q,h,k}$ is 1-Lipschitz, we deduce that

$$\mathcal{B}\left(x, \frac{\Delta_{\infty, K}}{2}\right) \subset \left\{y \mid d_{Q,k,h}(y) \geq \frac{\Delta_{\infty, K}}{2}\right\}.$$

Since $P\left(\mathcal{B}\left(x, \frac{\Delta_{\infty,K}}{2}\right)\right) \geq C(P)\left(\frac{\Delta_{\infty,K}}{2}\right)^{d'} \wedge 1$, Markov inequality yields that

$$\Delta_P^2 \geq C(P)\left(\frac{\Delta_{\infty,K}}{2}\right)^{d'+2} \wedge \frac{\Delta_{\infty,K}^2}{4}.$$

Thus we have

$$\sup_{x \in M} (d_{Q,h,k} - d_M)(x) = \Delta_{\infty,K} \leq C(P)^{-\frac{1}{d'+2}} \Delta_P^{\frac{2}{d'+2}} \vee 2\Delta_P.$$

Now, for $x \in \mathbb{R}^d$, we let $p \in M$ such that $\|x - p\| = d_M(x)$. Denote by $r = \|x - p\|$, and let t_j be such that $d_{Q,h,k}(p) = \sqrt{\|p - m(Q_{t_j,h})\|^2 + v(Q_{t_j,h})}$. Then

$$\begin{aligned} d_{Q,h,k}(x) &\leq \sqrt{\|x - m(Q_{t_j,h})\|^2 + v(Q_{t_j,h})} \\ &\leq \sqrt{d_{Q,h,k}^2(p) + r^2 + 2r\|p - m(Q_{t_j,h})\|} \\ &\leq \sqrt{d_{Q,h,k}^2(p) + r^2 + 2rd_{Q,h,k}(p)} \\ &= d_M(x) + (d_{Q,h,k}(p) - d_M(p)). \end{aligned}$$

Hence, $\sup_{x \in \mathbb{R}^d} (d_{Q,h,k} - d_M)(x) = \sup_{x \in M} (d_{Q,h,k} - d_M)(x) = \Delta_{\infty,K}$.

On the other hand, we have $d_{Q,h,k} \geq d_{Q,h}$ along with $\|d_{Q,h} - d_{P,h}\|_{\infty} \leq h^{-\frac{1}{2}}W_2(P, Q)$ (see, e.g., [CCSM11, Theorem 3.5]) as well as $d_{P,h} \geq d_M$. Hence

$$d_{Q,h,k} \geq d_M - h^{-\frac{1}{2}}W_2(P, Q).$$

4.4.2 Proofs for Section 4.2

Preliminary results about sub-Gaussian distributions

A sub-Gaussian measure P with variance $V^2 > 0$ is a measure $P \in \mathcal{P}(\mathbb{R}^d)$ such that $P(\mathcal{B}(0, t)^c) \leq \exp(-\frac{t^2}{2V^2})$ for all $t > V$. The set of such measures is denoted by $\mathcal{P}^{(V)}(\mathbb{R}^d)$. We derive some properties for sub-Gaussian distributions:

Lemma 4.36. *Let $P \in \mathcal{P}^{(V)}(\mathbb{R}^d)$, a sub-Gaussian measure with variance $V^2 > 0$. Then, the following bounds are satisfied:*

$$\begin{aligned} P\|u\| &\leq 3V, \\ P\|u\|^2 &\leq 3V^2, \\ P\|u\|^4 &\leq 9V^4. \end{aligned}$$

As a consequence, for all $n \in \mathbb{N}^*$, with $R = \sqrt{4V^2(1+p)\log(n)}$ we have:

$$\begin{aligned} P\mathbb{1}_{\|u\| \geq R} &\leq n^{-2p-2}, \\ P(\|u\|\mathbb{1}_{\|u\| \geq R}) &\leq \sqrt{3}Vn^{-(1+p)}, \\ P(\|u\|^2\mathbb{1}_{\|u\| \geq R}) &\leq 3V^2n^{-(1+p)}, \end{aligned}$$

and

$$\mathbb{P}\left(\sup_{i \in [1, n]} \|X_i\| \geq R\right) \leq n^{-2p-1},$$

where $\mathbb{X}_n = \{X_1, X_2, \dots, X_n\}$ is an n -sample from P .

For all $Q \in \mathcal{P}_h(P)$, that is, such that hQ is a sub-measure of P with P -mass h , then we have

$$Q\|u\|^2 \leq \frac{3V^2}{h}$$

and

$$\|Qu\| \leq \frac{\sqrt{3}V}{\sqrt{h}}.$$

Proof. For the first inequality, note that

$$\begin{aligned} P\|u\| &= \int_{u \in \mathbb{R}^d} \int_{t=0}^{+\infty} \mathbb{1}_{t \leq \|u\|} dt dP(u) \\ &= \int_{t=0}^V P(\|u\| \geq t) dt + \int_{t=V}^{+\infty} P(\|u\| \geq t) dt \\ &\leq V + \int_{t=V}^{+\infty} e^{-\frac{t^2}{2V^2}} dt \end{aligned}$$

moreover, for X distributed according to the standard normal distribution $\mathcal{N}(0, 1)$, we have

$$\begin{aligned} \int_{t=V}^{+\infty} \exp\left(-\frac{t^2}{2V^2}\right) dt &= \sqrt{2\pi}V \mathbb{P}(X \geq 1) \\ &\leq \sqrt{2\pi}V \frac{\mathbb{E}[\exp(\lambda X)]}{\exp(\lambda)} \\ &= \sqrt{2\pi}V \exp\left(-\lambda + \frac{\lambda^2}{2}\right). \end{aligned}$$

The choice $\lambda = 1$ yields

$$P\|u\| \leq \left(1 + \sqrt{2\pi} \exp\left(-\frac{1}{2}\right)\right) V.$$

For the second inequality, note that

$$\begin{aligned} P\|u\|^2 &= P\|u\|^2 \mathbb{1}_{\|u\| \leq V} + P\|u\|^2 \mathbb{1}_{\|u\| > V} \\ &\leq V^2 + P\|u\|^2 \mathbb{1}_{\|u\| > V}. \end{aligned}$$

Then,

$$\begin{aligned} P\|u\|^2 \mathbb{1}_{\|u\| > V} &= \int_{u \in \mathbb{R}^d} \mathbb{1}_{\|u\| > V} \int_{t=0}^{+\infty} 2t \mathbb{1}_{t \leq \|u\|} dt dP(u) \\ &= \int_{t=0}^V 2tP(\|u\| > V) dt + \int_{t=V}^{+\infty} 2tP(\|u\| \geq t) dt \\ &\leq \int_{t=0}^V 2te^{-\frac{1}{2}} dt + \int_{t=V}^{+\infty} 2te^{-\frac{t^2}{2V^2}} dt \\ &= 3V^2 e^{-\frac{1}{2}}. \end{aligned}$$

Thus,

$$P\|u\|^2 \mathbb{1}_{\|u\| > V} \leq 2V^2. \tag{4.21}$$

And,

$$P\|u\|^2 \leq 3V^2.$$

For the third inequality, note that

$$\begin{aligned} P\|u\|^4 &= P\|u\|^4 \mathbb{1}_{\|u\| \leq V} + P\|u\|^4 \mathbb{1}_{\|u\| > V} \\ &\leq V^4 + P\|u\|^4 \mathbb{1}_{\|u\| > V}. \end{aligned}$$

Then,

$$\begin{aligned} P\|u\|^4 \mathbb{1}_{\|u\| > V} &= \int_{u \in \mathbb{R}^d} \mathbb{1}_{\|u\| > V} \int_{t=0}^{+\infty} 4t^3 \mathbb{1}_{t \leq \|u\|} dt dP(u) \\ &= \int_{t=0}^V 4t^3 \mathbb{P}(\|u\| > V) dt + \int_{t=V}^{+\infty} 4t^3 \mathbb{P}(\|u\| > t) dt \\ &\leq \int_{t=0}^V 4t^3 e^{-\frac{1}{2}t^2} dt + \int_{t=V}^{+\infty} 4t^3 e^{-\frac{t^2}{2V^2}} dt \\ &= V^4 e^{-\frac{1}{2}} + V^4 \int_{u=1}^{+\infty} 4u^3 e^{-\frac{u^2}{2}} du \end{aligned}$$

After an integration by parts,

$$\int_{u=1}^{+\infty} 4u^3 e^{-\frac{u^2}{2}} du = \left[-4u^2 e^{-\frac{u^2}{2}} \right]_{u=1}^{+\infty} + 8 \int_{u=1}^{+\infty} u e^{-\frac{u^2}{2}} du = 12e^{-\frac{1}{2}}.$$

As a consequence,

$$P\|u\|^4 \leq 9V^4.$$

The two following inequalities come from Cauchy-Schwarz inequality and the definition of a sub-Gaussian measure. The fifth inequality holds since

$$\mathbb{P} \left(\sup_{i \in [1, n]} \|X_i\| \geq R \right) \leq n \mathbb{P}(\|X_1\| \geq R) \leq n^{-2p-1}.$$

Let $Q \in \mathcal{P}_h(P)$. We may write

$$\begin{aligned} Q\|u\|^2 &\leq \frac{1}{h} P\|u\|^2 \\ &\leq \frac{1}{h} [P\|u\|^2 \mathbb{1}_{\|u\| \leq V} + P\|u\|^2 \mathbb{1}_{\|u\| > V}] \\ &\leq \frac{V^2}{h} + \frac{P\|u\|^2 \mathbb{1}_{\|u\| > V}}{h} \\ &\leq 3 \frac{V^2}{h}. \end{aligned}$$

The last inequality comes from (4.21). □

Lemma 4.37. *If Y is a random variable sampled from a distribution P in $\mathcal{P}^K(\mathbb{R}^d)$ and Z is independent from Y and sampled from a distribution Q' in $\mathcal{P}^{(\sigma)}(\mathbb{R}^d)$ for some $\sigma > 0$. Then, the distribution Q of the random variable $X = Y + Z$ is sub-Gaussian with variance $V^2 = (K + \sigma)^2$, that is in $\mathcal{P}^{(K+\sigma)}(\mathbb{R}^d)$.*

Moreover,

$$W_1(P, Q) \leq 3\sigma \text{ and } W_2(P, Q) \leq \sqrt{3}\sigma.$$

Proof. For all $z \geq \sigma$ it holds

$$\begin{aligned} \mathbb{P}_{(Y,Z)}(\|Y + Z\| \geq K + z) &\leq \mathbb{P}_Z(\|Z\| \geq z) \\ &\leq \exp -\frac{z^2}{2\sigma^2} \\ &\leq \exp -\frac{(z + K)^2}{2(\sigma + K)^2}. \end{aligned}$$

The L_1 -Wasserstein distance between P and Q satisfies

$$W_1(P, Q) = \mathbb{E}_{(Y,Z)} [\|(Y + Z) - Y\|] = Q'\|u\|,$$

which is bounded from above by 3σ according to Lemma 4.36.

The L_2 -Wasserstein distance between P and Q satisfies

$$W_2(P, Q) = \sqrt{\mathbb{E}_{(Y,Z)} [\|(Y + Z) - Y\|^2]} = \sqrt{Q'\|u\|^2},$$

which is bounded from above by $\sqrt{3}\sigma$ according to Lemma 4.36. □

Proof of Proposition 4.15

For every $t = (t_1, t_2, \dots, t_k) \in \mathbb{R}^{d(k)}$, we note $c_i = \frac{\sum_{X \in \mathcal{C}(t_i)} X}{|\mathcal{C}(t_i)|}$. Then,

$$\begin{aligned} &Q_n \min_{i \in [1, k]} \|u - m(Q_{n t_i, h})\|^2 + v(Q_{n t_i, h}) \\ &= \sum_{i=1}^k \frac{1}{n} \sum_{X \in \mathcal{C}(t_i)} \|X - m(Q_{n t_i, h})\|^2 + v(Q_{n t_i, h}) \\ &= \sum_{i=1}^k \frac{1}{n} \sum_{X \in \mathcal{C}(t_i)} \|X\|^2 - 2\langle X - t_i, m(Q_{n t_i, h}) \rangle + (d_{Q_{n, h}}^2(t_i) - \|t_i\|^2) \\ &= \frac{1}{n} \sum_{j=1}^n \|X_j\|^2 + \sum_{i=1}^k \frac{|\mathcal{C}(t_i)|}{n} (-2\langle c_i - t_i, m(Q_{n t_i, h}) \rangle + (d_{Q_{n, h}}^2(t_i) - \|t_i\|^2)) \\ &\geq \frac{1}{n} \sum_{j=1}^n \|X_j\|^2 + \sum_{i=1}^k \frac{|\mathcal{C}(t_i)|}{n} (d_{Q_{n, h}}^2(c_i) - \|c_i\|^2) \\ &= \sum_{i=1}^k \frac{1}{n} \sum_{X \in \mathcal{C}(t_i)} \|X - m(Q_{n c_i, h})\|^2 + v(Q_{n c_i, h}) \\ &\geq \sum_{i=1}^k \frac{1}{n} \sum_{X \in \mathcal{C}(c_i)} \|X - m(Q_{n c_i, h})\|^2 + v(Q_{n c_i, h}) \\ &= Q_n \min_{i \in [1, k]} \|u - m(Q_{n c_i, h})\|^2 + v(Q_{n c_i, h}). \end{aligned}$$

We used the semi-concavity property of the distance-to-measure, see Lemma 4.28. The algorithm converges since the number of possible decomposition of \mathbb{X}_n in Voronoi cells is finite, and the loss decreases whenever $Q_{n t_i, h} \notin \mathcal{P}_{c_i, h}(Q_n)$ according to Lemma 4.28 or when the Voronoi cells associated to the t_i s and the c_i s are different.

Proof of Theorem 4.16

Let γ and $\hat{\gamma}$ the functions defined for $(t, x) \in \mathbb{R}^{d(k)} \times \mathbb{R}^d$ with $t = (t_1, t_2, \dots, t_k)$, by:

$$\gamma(t, x) = \min_{i \in [1, k]} -2\langle x, m(Q_{t_i, h}) \rangle + \|m(Q_{t_i, h})\|^2 + v(Q_{t_i, h}),$$

and

$$\hat{\gamma}(t, x) = \min_{i \in [1, k]} -2\langle x, m(Q_{nt_i, h}) \rangle + \|m(Q_{nt_i, h})\|^2 + v(Q_{nt_i, h}).$$

According to Lemma 4.37, $Q \in \mathcal{P}^{(V)}(\mathbb{R}^d)$ with $V = \sigma + K$. The proof of Theorem 4.16 is based on the two following deviation Lemmas.

Lemma 4.38. *If Q is sub-Gaussian with variance V^2 , then, for every $p > 0$, with probability larger than $1 - 2n^{-p}$, we have*

$$\sup_{t \in \mathbb{R}^{d(k)}} |(Q - Q_n)\gamma(t, u)| \leq C \frac{V^2 \sqrt{kd}(1+p)^{\frac{3}{2}} \log(n)^{\frac{3}{2}}}{h\sqrt{n}},$$

for some absolute positive constant C .

The proof of Lemma 4.38 is deferred to Section 4.4.2.

Lemma 4.39. *Assume that Q is sub-Gaussian with variance V^2 , then, for every $p > 0$, with probability larger than $1 - 9n^{-p}$, we have*

$$\sup_{t \in \mathbb{R}_d^{(k)}} |Q_n(\gamma - \hat{\gamma})(t, u)| \leq CV^2 \frac{\sqrt{d}(p+1)^{\frac{3}{2}} \log(n)^{\frac{3}{2}}}{h\sqrt{n}}.$$

As well, the proof of Lemma 4.39 is deferred to Section 4.4.2.

We are now in position to prove Theorem 4.16.

Let

$$s = \arg \min \left\{ Q\gamma(t, u) \mid t = (t_1, t_2, \dots, t_k) \in \mathbb{R}^{d(k)} \right\},$$

$$\hat{s} = \arg \min \left\{ Q_n \hat{\gamma}(t, u) \mid t = (t_1, t_2, \dots, t_k) \in \mathbb{R}^{d(k)} \right\}$$

and

$$\tilde{s} = \arg \min \left\{ Q_n \gamma(t, u) \mid t = (t_1, t_2, \dots, t_k) \in \mathbb{R}^{d(k)} \right\}.$$

With these notations, for all $x \in \mathbb{R}^d$, $d_{Q, h, k}^2(x) = \|x\|^2 + \gamma(s, x)$ and $d_{Q_n, h, k}^2(x) = \|x\|^2 + \hat{\gamma}(\hat{s}, x)$. We intend to bound $l(s, \hat{s}) = Q(d_{Q_n, h, k}^2(u) - d_{Q, h, k}^2(u))$, which is also equal to $l(s, \hat{s}) = Q(\gamma(\hat{s}, u) - Q\gamma(s, u))$.

We have that:

$$\begin{aligned} l(s, \hat{s}) &= Q\gamma(\hat{s}, u) - Q_n\gamma(\hat{s}, u) + Q_n\gamma(\hat{s}, u) - Q_n\gamma(\tilde{s}, u) + Q_n\gamma(\tilde{s}, u) - Q\gamma(s, u) \\ &\leq \sup_{t \in \mathbb{R}^{d(k)}} (Q - Q_n)\gamma(t, u) + Q_n(\gamma - \hat{\gamma})(\hat{s}, u) \\ &\quad + Q_n(\hat{\gamma}(\hat{s}, u) - \hat{\gamma}(\tilde{s}, u)) + Q_n(\hat{\gamma} - \gamma)(\tilde{s}, u) + \sup_{t \in \mathbb{R}^{d(k)}} (Q_n - Q)\gamma(t, u), \end{aligned}$$

where we used the fact that $Q_n\gamma(\tilde{s}, u) \leq Q_n\gamma(s, u)$. Now, since $Q_n(\hat{\gamma}(\hat{s}, u) - \hat{\gamma}(\tilde{s}, u)) \leq 0$, we get:

$$\begin{aligned} l(s, \hat{s}) &\leq \sup_{t \in \mathbb{R}^{d(k)}} (Q - Q_n)\gamma(t, u) + \sup_{t \in \mathbb{R}^{d(k)}} (Q_n - Q)\gamma(t, u) \\ &\quad + \sup_{t \in \mathbb{R}^{d(k)}} Q_n(\gamma - \hat{\gamma})(t, u) + \sup_{t \in \mathbb{R}^{d(k)}} Q_n(\hat{\gamma} - \gamma)(t, u). \end{aligned}$$

Combining Lemma 4.38 and Lemma 4.39 entails, with probability larger than $1 - 10n^{-p}$,

$$l(s, \hat{s}) \leq CV^2 \sqrt{kd} \frac{(p+1)^{\frac{3}{2}} \log(n)^{\frac{3}{2}}}{h\sqrt{n}}.$$

It remains to bound $|Pd_{Q_n, h, k}^2 - Qd_{Q_n, h, k}^2|$ as well as $|Pd_{Q, h, k}^2 - Qd_{Q, h, k}^2|$. To this aim we recall that $X = Y + Z$, Z being sub-Gaussian with variance σ^2 . Thus, denoting by $s_j(x) = \arg \min_{j \in [1, k]} \|x - m(Q_{s_j, h})\|^2 + v(Q_{s_j, h})$,

$$\begin{aligned} Pd_{Q, h, k}^2 - Qd_{Q, h, k}^2 &\leq \mathbb{E}_{(Y, Z)} [\|Y - m(Q_{s_j(Y+Z), h})\|^2 + v(Q_{s_j(Y+Z), h}) \\ &\quad - (\|Y + Z - m(Q_{s_j(Y+Z), h})\|^2 + v(Q_{s_j(Y+Z), h}))] \\ &\leq \mathbb{E}_Z \|Z\|^2 + 2\mathbb{E}_{(Y, Z)} \max_{j \in [1, k]} |\langle Z, m(Q_{s_j, h}) - Y \rangle| \\ &\leq 3\sigma^2 + 2\sqrt{3}\sigma \left(\max_{j \in [1, k]} \|m(Q_{s_j, h})\| + K \right) \\ &\leq \frac{C\sigma K}{\sqrt{h}}, \end{aligned}$$

using Cauchy-Schwarz inequality, Lemma 4.36 and $\sigma \leq K$. The converse bound on $Qd_{Q, h, k}^2 - Pd_{Q, h, k}^2$ may be proved the same way. Similarly, we may write

$$\begin{aligned} Pd_{Q_n, h, k}^2 - Qd_{Q_n, h, k}^2 &\leq 3\sigma^2 + 2\sqrt{3}\sigma \left(\max_{j \in [1, k]} \|m(Q_n s_j, h)\| + K \right) \\ &\leq 3\sigma^2 + 2\sqrt{3}\sigma \left(\max_{j \in [1, k]} \|m(Q_{s_j, h})\| + \sup_{t \in \mathbb{R}^d} \|m(Q_{t, h}) - m(Q_n t, h)\| + K \right) \\ &\leq 3\sigma^2 + 2\sqrt{3}\sigma \left(\max_{j \in [1, k]} \|m(Q_{s_j, h})\| + C(K + \sigma)\sqrt{d} \frac{(p+1) \log(n)}{h\sqrt{n}} + K \right) \\ &\leq \frac{C\sigma K}{\sqrt{h}} + \frac{C\sigma K \sqrt{d}(p+1) \log(n)}{h\sqrt{n}}, \end{aligned}$$

according to Lemma 4.36 and (4.24) in the proof of Lemma 4.39. The bound on $Qd_{Q_n, h, k}^2 - Pd_{Q_n, h, k}^2$ derives from the same argument. Collecting all pieces, we have

$$\begin{aligned} |P(d_{Q_n, h, k}^2 - d_{Q, h, k}^2)| &\leq |Q(d_{Q_n, h, k}^2 - d_{Q, h, k}^2)| + \frac{C\sigma K \sqrt{d}(p+1) \log(n)}{h\sqrt{n}} + \frac{C\sigma K}{\sqrt{h}} \\ &\leq \frac{C\sigma K \sqrt{d}(p+1) \log(n)}{h\sqrt{n}} + \frac{CkK^2 \sqrt{dk}((p+1) \log(n))^{\frac{3}{2}}}{h\sqrt{n}} + \frac{C\sigma K}{\sqrt{h}}, \end{aligned}$$

where we used $\sigma \leq K$.

Proof of Lemma 4.38

With the notation $l_{t_i}(x) = -2\langle x, m(Q_{t_i, h}) \rangle + \|m(Q_{t_i, h})\|^2 + v(Q_{t_i, h})$, we get that:

$$\sup_{t \in \mathbb{R}^d (k)} |(Q - Q_n)\gamma(t, u)| = \sup_{t \in \mathbb{R}^d (k)} \left| (Q - Q_n) \min_{i \in [1, k]} l_{t_i}(u) \right|.$$

First we note that since Q is sub-Gaussian with variance V^2 , we have from Lemma 4.36, for every $c \in \mathbb{R}^d$,

$$\|m(Q_{c, h})\|^2 + v(Q_{c, h}) = Q_{c, h} \|u\|^2 \leq \frac{3V^2}{h}. \tag{4.22}$$

Set $R = 2V\sqrt{\log(n) + \lambda}$ and $\lambda = p \log(n)$. Then, according to Lemma 4.36, with probability larger than $1 - n^{-2p-1}$,

$$\max_{i \in \llbracket 1, n \rrbracket} \|X_i\| \leq R. \quad (4.23)$$

We may then write

$$\begin{aligned} \sup_{t \in \mathbb{R}^{d(k)}} |(Q - Q_n)\gamma(t, u)| &= \sup_{t \in \mathbb{R}^{d(k)}} \left| \frac{1}{n} \sum_{i=1}^n \gamma(t, X_i) - Q\gamma(t, u) \right| \\ &\leq \sup_{t \in \mathbb{R}^{d(k)}} \left| \frac{1}{n} \sum_{i=1}^n \gamma(t, X_i) \mathbb{1}_{\|X_i\| \leq R} - Q\gamma(t, u) \mathbb{1}_{\|u\| \leq R} \right| \\ &\quad + \sup_{t \in \mathbb{R}^{d(k)}} Q|\gamma(t, u) \mathbb{1}_{\|u\| > R}| + \sup_{t \in \mathbb{R}^{d(k)}} \left| \frac{1}{n} \sum_{i=1}^n \gamma(t, X_i) \mathbb{1}_{\|X_i\| > R} \right|. \end{aligned}$$

According to (4.23), the last part is 0 with probability larger than $1 - n^{-2p-1}$. Moreover, according to Lemma 4.36,

$$\begin{aligned} Q|\gamma(t, u) \mathbb{1}_{\|u\| > R}| &\leq Q \mathbb{1}_{\|u\| > R} \sup_{i \in \llbracket 1, k \rrbracket} 2|\langle u, m(Q_{t_i, h}) \rangle| + \|m(Q_{t_i, h})\|^2 + v(Q_{t_i, h}) \\ &\leq 2\sqrt{3} \frac{V}{\sqrt{h}} Q \|u\| \mathbb{1}_{\|u\| > R} + \frac{3V^2}{h} Q \mathbb{1}_{\|u\| > R} \\ &\leq \frac{3V^2}{h} n^{-1-p} (2\sqrt{h} + n^{-1-p}) \\ &\leq \frac{9V^2}{h} n^{-1-p}. \end{aligned}$$

It remains to bound

$$\sup_{t \in \mathbb{R}^{d(k)}} |(Q - Q_n)\gamma(t, u) \mathbb{1}_{\|u\| \leq R}|.$$

Since from Lemma 4.36, for every t and u , $|\gamma(t, u) \mathbb{1}_{\|u\| \leq R}| \leq \left(R + \frac{V\sqrt{3}}{\sqrt{h}}\right)^2 := Z$, then Corollary 3.44 entails

$$\mathbb{P} \left(\sup_{t \in \mathbb{R}^{d(k)}} |(Q - Q_n)\gamma(t, u) \mathbb{1}_{\|u\| \leq R}| \geq \mathbb{E} \sup_{t \in \mathbb{R}^{d(k)}} |(Q - Q_n)\gamma(t, u) \mathbb{1}_{\|u\| \leq R}| + Z \sqrt{\frac{2\lambda}{n}} \right) \leq e^{-\lambda} = n^{-p}.$$

Then, Theorem 3.55 yields

$$\mathbb{E} \sup_{t \in \mathbb{R}^{d(k)}} |(Q - Q_n)\gamma(t, u) \mathbb{1}_{\|u\| \leq R}| \leq 24 \frac{Z}{\sqrt{n}} \mathbb{E} \left[\int_0^{\frac{1}{2}} \sqrt{\log \left(N'_{\|\cdot\|} \left(\frac{u}{2}, \frac{1}{Z\sqrt{n}} (\mathcal{F}^{(k)} \cup -\mathcal{F}^{(k)})(X_1^n) \right) \right)} du \right],$$

with $\frac{1}{Z\sqrt{n}} \mathcal{F}^{(k)} = \left\{ \frac{1}{Z\sqrt{n}} \mathbb{1}_{\|u\| \leq R} \min_{i \in \llbracket 1, k \rrbracket} l_{t_i}(u) \right\}$. According to (3.39), (3.40) and (3.37), it holds that

$$\begin{aligned} N'_{\|\cdot\|} \left(\frac{u}{2}, \frac{1}{Z\sqrt{n}} (\mathcal{F}^{(k)} \cup -\mathcal{F}^{(k)})(X_1^n) \right) &\leq 2N'_{\|\cdot\|} \left(\frac{u}{2}, \frac{1}{Z\sqrt{n}} (\mathcal{F}^{(k)})(X_1^n) \right) \\ &\leq 2 \left(N'_{\|\cdot\|} \left(\frac{u}{2\sqrt{k}}, \frac{1}{Z\sqrt{n}} (\mathcal{F})(X_1^n) \right) \right)^k \\ &\leq 2 \left(N'_{\|\cdot\|} \left(\frac{u}{4\sqrt{k}}, \frac{1}{Z\sqrt{n}} (\mathcal{G}_1)(X_1^n) \right) \times N'_{\|\cdot\|} \left(\frac{u}{4\sqrt{k}}, \frac{1}{Z\sqrt{n}} (\mathcal{G}_2)(X_1^n) \right) \right)^k \end{aligned}$$

with $\mathcal{G}_1 = \left\{ x \mapsto \frac{-2\langle x, m(Q_{t,h}) \rangle \mathbb{1}_{\|x\| \leq R}}{Z\sqrt{n}} \mid t \in \mathbb{R}^d \right\}$ and $\mathcal{G}_2 = \left\{ x \mapsto \frac{(\|m(Q_{t,h})\|^2 + v(Q_{t,h})) \mathbb{1}_{\|x\| \leq R}}{Z\sqrt{n}} \mid t \in \mathbb{R}^d \right\}$. According to (3.45) and Lemma 3.53,

$$N'_{\|\cdot\|}(\delta, \mathcal{G}_1(X_1^n)) \leq \left(\frac{2}{\delta}\right)^{V'(d+2)},$$

for some absolute positive constant V' . Also, $N'_{\|\cdot\|}(\delta, \mathcal{G}_2(X_1^n)) \leq N'_{\|\cdot\|}(\delta, \mathcal{G}_3(X_1^n)) \leq \frac{2}{\delta}$ with $\mathcal{G}_3 = \left\{ x \mapsto \frac{t \mathbb{1}_{\|x\| \leq R}}{\sqrt{n}} \mid t \in [0, 1] \right\}$.

As a consequence, with probability larger than $1 - 2n^{-p}$, it holds

$$\begin{aligned} \sup_{t \in \mathbb{R}^{d(k)}} |(Q - Q_n)\gamma(t, u)| &\leq CZ \frac{\sqrt{k}\sqrt{d+2}}{\sqrt{n}} + \frac{9V^2}{h} n^{-1-p} + Z\sqrt{\frac{2p \log(n)}{n}} \\ &\leq C \frac{V^2 \sqrt{k}d(1+p)^{\frac{3}{2}} \log(n)^{\frac{3}{2}}}{h\sqrt{n}}, \end{aligned}$$

for some positive absolute constant C .

Proof of Lemma 4.39

For $t \in \mathbb{R}^{d(k)}$, we get that:

$$\begin{aligned} |\gamma(t, x) - \hat{\gamma}(t, x)| &\leq \max_{j \in [1, k]} | -2\langle x, m(Q_{t_j, h}) - m(Q_{n t_j, h}) \rangle + (M(Q_{t_j, h}) - M(Q_{n t_j, h})) | \\ &\leq 2\|x\| \max_{j \in [1, k]} \|m(Q_{t_j, h}) - m(Q_{n t_j, h})\| + \max_{j \in [1, k]} |M(Q_{t_j, h}) - M(Q_{n t_j, h})|. \end{aligned}$$

Let $t \in \mathbb{R}^d$, and denote by $r = \delta_{Q, h}(t)$, $r_n = \delta_{Q_n, h}(t)$, and $R = 2V\sqrt{(p+1)\log(n)}$. We may write for $Q_{t, h} = \frac{1}{h}Q\mathbb{1}_{\mathcal{B}(t, r)} + \frac{\alpha}{h}Q\mathbb{1}_{\partial\mathcal{B}(t, r)}$ and for $Q_{n t, h} = \frac{1}{h}Q_n\mathbb{1}_{\mathcal{B}(t, r_n)} + \frac{\alpha_n}{h}Q_n\mathbb{1}_{\partial\mathcal{B}(t, r_n)}$

$$\begin{aligned} \|m(Q_{t, h}) - m(Q_{n t, h})\| &\leq \frac{1}{h} (\|Qu\mathbb{1}_{\mathcal{B}(t, r)}(u) + \alpha Qu\mathbb{1}_{\partial\mathcal{B}(t, r)}(u) - Qu\mathbb{1}_{\mathcal{B}(t, r_n)}(u) - \alpha_n Qu\mathbb{1}_{\partial\mathcal{B}(t, r_n)}(u)\| \\ &\quad + \|Qu\mathbb{1}_{\mathcal{B}(t, r_n)}(u) - Q_n u\mathbb{1}_{\mathcal{B}(t, r_n)}(u)\| + \alpha_n \|Qu\mathbb{1}_{\partial\mathcal{B}(t, r_n)}(u) - Q_n u\mathbb{1}_{\partial\mathcal{B}(t, r_n)}(u)\|) \\ &\leq \frac{1}{h} (Q\|u\| |\mathbb{1}_{\mathcal{B}(t, r)} + \alpha\mathbb{1}_{\partial\mathcal{B}(t, r)} - \mathbb{1}_{\mathcal{B}(t, r_n)} - \alpha_n\mathbb{1}_{\partial\mathcal{B}(t, r_n)}|(u) + \|(Q - Q_n)u\mathbb{1}_{\mathcal{B}(t, r_n)}(u)\| \\ &\quad + \|(Q - Q_n)u\mathbb{1}_{\partial\mathcal{B}(t, r_n)}(u)\|) \\ &\leq \frac{1}{h} (RQ |\mathbb{1}_{\mathcal{B}(t, r)} + \alpha\mathbb{1}_{\partial\mathcal{B}(t, r)} - \mathbb{1}_{\mathcal{B}(t, r_n)} - \alpha_n\mathbb{1}_{\partial\mathcal{B}(t, r_n)}|(u) + \|(Q - Q_n)u\mathbb{1}_{\mathcal{B}(t, r_n)}(u)\| \\ &\quad + \|(Q - Q_n)u\mathbb{1}_{\partial\mathcal{B}(t, r_n)}(u)\| + Q\|u\| \mathbb{1}_{\|u\| > R}). \end{aligned}$$

Moreover, considering the case $r > r_n$ for instance, we have:

$$\begin{aligned} &Q |\mathbb{1}_{\mathcal{B}(t, r)} + \alpha\mathbb{1}_{\partial\mathcal{B}(t, r)} - \mathbb{1}_{\mathcal{B}(t, r_n)} - \alpha_n\mathbb{1}_{\partial\mathcal{B}(t, r_n)}|(u) \\ &= Q(\mathcal{B}(t, r)) + \alpha Q(\partial\mathcal{B}(t, r)) - Q(\mathcal{B}(t, r_n)) - \alpha_n Q(\partial\mathcal{B}(t, r_n)) \\ &= h - Q(\mathcal{B}(t, r_n)) - \alpha_n Q(\partial\mathcal{B}(t, r_n)) \\ &\leq |(Q_n - Q)\mathbb{1}_{\mathcal{B}(t, r_n)} + \alpha_n\mathbb{1}_{\partial\mathcal{B}(t, r_n)}|(u) \\ &\leq |Q_n - Q| \mathbb{1}_{\mathcal{B}(t, r_n)}(u) + |Q_n - Q| \mathbb{1}_{\partial\mathcal{B}(t, r_n)}(u). \end{aligned}$$

The same inequality holds when $r \leq r_n$.

On the event described in Lemma 3.58, we have that

$$\begin{aligned} \|(Q - Q_n)u\mathbb{1}_{\mathcal{B}(t,r_n)}(u)\| &\leq CV\sqrt{d}\frac{(p+1)\log(n)}{\sqrt{n}}, \\ \|(Q - Q_n)u\mathbb{1}_{\partial\mathcal{B}(t,r_n)}(u)\| &\leq CV\sqrt{d}\frac{(p+1)\log(n)}{\sqrt{n}}, \\ |Q_n(\mathcal{B}(t,r_n)) - Q(\mathcal{B}(t,r_n))| &\leq C\sqrt{d}\frac{\sqrt{(p+1)\log(n)}}{\sqrt{n}}, \\ |Q_n(\partial\mathcal{B}(t,r_n)) - Q(\partial\mathcal{B}(t,r_n))| &\leq C\sqrt{d}\frac{\sqrt{(p+1)\log(n)}}{\sqrt{n}}, \\ Q\|u\|\mathbb{1}_{\|u\|>R} &\leq 2Vn^{-(p+1)}, \end{aligned}$$

from Lemma 4.36. Thus,

$$\sup_{t \in \mathbb{R}^d} \|m(Q_{t,h}) - m(Q_{n,t,h})\| \leq \frac{CV\sqrt{d}(p+1)\log(n)}{h\sqrt{n}}. \quad (4.24)$$

As well, $\sup_{t \in \mathbb{R}^d} |M(Q_{t,h}) - M(Q_{n,t,h})|$ is bounded from above by

$$\begin{aligned} &\frac{1}{h}(R^2|Q_n - Q|\mathbb{1}_{\mathcal{B}(t,r_n)}(u) + R^2|Q_n - Q|\mathbb{1}_{\partial\mathcal{B}(t,r_n)}(u) + \|(Q - Q_n)u\mathbb{1}_{\mathcal{B}(t,r_n)}(u)\| \\ &\quad + \|(Q - Q_n)u\mathbb{1}_{\partial\mathcal{B}(t,r_n)}(u)\| + Q\|u\|^2\mathbb{1}_{\|u\|>R}). \end{aligned}$$

Using Lemma 3.58 again, we get

$$\begin{aligned} |(Q - Q_n)\|u\|^2\mathbb{1}_{\mathcal{B}(t,r_n)}| &\leq CV^2\sqrt{d}\frac{(p+1)^{\frac{3}{2}}\log(n)^{\frac{3}{2}}}{\sqrt{n}} \\ |(Q - Q_n)\|u\|^2\mathbb{1}_{\partial\mathcal{B}(t,r_n)}| &\leq CV^2\sqrt{d}\frac{(p+1)^{\frac{3}{2}}\log(n)^{\frac{3}{2}}}{\sqrt{n}} \\ |(Q - Q_n)\mathbb{1}_{\mathcal{B}(t,r_n)}| &\leq C\sqrt{d}\frac{\sqrt{(p+1)\log(n)}}{\sqrt{n}} \\ |(Q - Q_n)\mathbb{1}_{\partial\mathcal{B}(t,r_n)}| &\leq C\sqrt{d}\frac{\sqrt{(p+1)\log(n)}}{\sqrt{n}} \\ Q\|u\|^2\mathbb{1}_{\|u\|>R} &\leq 3V^2n^{-(p+1)}, \end{aligned}$$

from Lemma 4.36. Collecting all pieces leads to

$$|\gamma(t,x) - \hat{\gamma}(t,x)| \leq C\|x\|\frac{V\sqrt{d}(p+1)\log(n)}{h\sqrt{n}} + CV^2\frac{\sqrt{d}(p+1)^{\frac{3}{2}}\log(n)^{\frac{3}{2}}}{h\sqrt{n}}. \quad (4.25)$$

At last, from Lemma 4.36

$$\mathbb{P}\left\{\max_i \|X_i\| \geq R\right\} \leq n^{-2p-1},$$

we deduce that

$$Q_n |\gamma(t,u) - \hat{\gamma}(t,u)| \leq CV^2\frac{\sqrt{d}(p+1)^{\frac{3}{2}}\log(n)^{\frac{3}{2}}}{h\sqrt{n}},$$

with probability larger than $1 - 9n^{-p}$.

Proof of Proposition 4.17

Putting bounds obtained in Theorem 4.16 and Proposition 4.13 together yields:

$$\begin{aligned}
 |Pd_{Q_n, h, k}^2(u) - d_{P, h}^2(u)| &\leq C\sqrt{kd} \frac{K^2((p+1)\log(n))^{\frac{3}{2}}}{h\sqrt{n}} + C \frac{K\sigma}{\sqrt{h}} \\
 &\quad + 3\|d_{Q, h}^2 - d_{P, h}^2\|_{\infty, \mathcal{B}(0, K)} + Pd_{P, h, k}^2(u) - d_{P, h}^2(u) + 4W_1(P, Q) \sup_{s \in \mathbb{R}^d} \|m(P_{s, h})\| \\
 &\leq C\sqrt{kd} \frac{K^2((p+1)\log(n))^{\frac{3}{2}}}{h\sqrt{n}} + C \frac{K\sigma}{\sqrt{h}} \\
 &\quad + 3\|d_{Q, h}^2 - d_{P, h}^2\|_{\infty, \mathcal{B}(0, K)} + C_P k^{-\frac{2}{d}} + 4W_1(P, Q) \sup_{s \in \mathbb{R}^d} \|m(P_{s, h})\| \\
 &\leq C\sqrt{kd} \frac{K^2((p+1)\log(n))^{\frac{3}{2}}}{h\sqrt{n}} + C \frac{K\sigma}{\sqrt{h}} \\
 &\quad + 3 \frac{\sqrt{3}\sigma(3K + \sqrt{3}(K + \sigma))}{h} + C_P k^{-\frac{2}{d}} + 12\sqrt{3} \frac{K + \sigma}{\sqrt{h}} \sigma
 \end{aligned}$$

where we used Corollary 4.12 and the bound

$$\begin{aligned}
 \|d_{Q, h}^2 - d_{P, h}^2\|_{\infty, \mathcal{B}(0, K)} &\leq \|d_{Q, h} - d_{P, h}\|_{\infty, \mathcal{B}(0, K)} (\|d_{Q, h}\|_{\infty, \mathcal{B}(0, K)} + \|d_{P, h}\|_{\infty, \mathcal{B}(0, K)}) \\
 &\leq \frac{W_2(P, Q)}{\sqrt{h}} \left(\sup_{x \in \mathcal{B}(0, K)} \sqrt{\|x - m(Q_{x, h})\|^2 + v(Q_{x, h})} + 2K \right) \\
 &\leq \frac{\sqrt{3}\sigma}{\sqrt{h}} \left(\sqrt{K^2 + 2K\sqrt{3} \frac{\sigma + K}{\sqrt{h}}} + 3 \frac{(\sigma + K)^2}{h} + 2K \right) \\
 &\leq \frac{\sqrt{3}\sigma(3K + \sqrt{3}(K + \sigma))}{h}.
 \end{aligned}$$

These inequalities come from Proposition 2.2 and the inequalities in Lemma 4.37 and Lemma 4.36:

$$\begin{aligned}
 W_1(P, Q) &\leq 3\sigma \\
 W_2(P, Q) &\leq \sqrt{3}\sigma \\
 \sup_{x \in \mathbb{R}^d} \|m(Q_{x, h})\| &\leq \sqrt{3} \frac{V}{\sqrt{h}} \\
 \sup_{x \in \mathbb{R}^d} M(Q_{x, h}) &\leq 3 \frac{V^2}{h}.
 \end{aligned}$$

4.4.3 Proofs for Section 4.3

Proof of Lemma 4.20

Set $P \in \mathcal{P}^K(\mathbb{R}^d)$ for some $K > 0$, a distribution which puts no mass on boundaries of balls and half-spaces of P -mass h . Let $(x_n)_{n \in \mathbb{N}}$ be a sequence in $\overline{\mathbb{R}^d}$ such that the sequence $(m(P_{x_n, h}))_{n \in \mathbb{N}}$ converges to some element $l \in \mathbb{R}^d$. We intend to prove that $l \in \overline{\mathcal{M}}_h(P)$.

Since $\overline{\mathbb{R}^d}$ is compact, set x the limit in $\overline{\mathbb{R}^d}$ of a subsequence of $(x_n)_{n \in \mathbb{N}}$.

First, we assume that $P(\partial H(v, c_{P, h}(v))) = 0$ and $P(\partial \mathcal{B}(y, \delta_{P, h}(y))) = 0$ for all $y \in \mathbb{R}^d$ and $v \in S(0, 1)$. If $x \in \mathbb{R}^d$, then for n large enough, $x_n \in \mathbb{R}^d$. Moreover, $|\delta_{P, h}(x_n) - \delta_{P, h}(x)| \leq$

$\|x - x_n\| \rightarrow 0$. Thus, for all $y \notin \partial\mathcal{B}(x, \delta_{P,h}(x))$, $\mathbb{1}_{\mathcal{B}(x_n, \delta_{P,h}(x_n))}(y) \rightarrow \mathbb{1}_{\mathcal{B}(x, \delta_{P,h}(x))}(y)$. Since $P(\partial\mathcal{B}(x, \delta_{P,h}(x))) = 0$, this holds for P -almost all y . The dominated convergence theorem yields $m(P_{x_n, h}) \rightarrow m(P_{x, h})$.

If $x = v_\infty \in \overline{\mathbb{R}^d} \setminus \mathbb{R}^d$ and $x_n = v_\infty^n \in \overline{\mathbb{R}^d} \setminus \mathbb{R}^d$ for an infinite set of n . Up to a subsequence $c_{P,h}(v_\infty^n) \rightarrow c_{P,h}(v)$. Since $P(\partial H(v, c_{P,h}(v))) = 0$, for P -almost all y , $\mathbb{1}_{H(v_\infty^n, c_{P,h}(v_\infty^n))}(y) \rightarrow \mathbb{1}_{H(v, c_{P,h}(v))}(y)$. Again the dominated convergence theorem yields $m(P_{v_\infty^n, h}) \rightarrow m(P_{v_\infty, h})$.

Finally, if $x = v_\infty \in \overline{\mathbb{R}^d} \setminus \mathbb{R}^d$ and $x_n \in \mathbb{R}^d$ for n large enough, then Lemma 4.19 and the dominated convergence lemma yield $m(P_{x_n, h}) \rightarrow m(P_{v_\infty, h})$.

Thus, the limit l belongs to $\overline{\mathcal{M}_h(P)}$ which is closed. Since $\overline{\mathcal{M}_h(P)} \subset \mathcal{B}(0, K)$, this set is compact and the Carathéodory theorem in the finite space \mathbb{R}^d yields that its convex hull is also compact.

If $P \in \mathcal{P}_n(\mathbb{R}^d)$, then as noticed in the proof of Lemma 4.23, $\text{Conv}(\overline{\mathcal{M}_h(P_n)}) = \text{Conv}(\hat{\mathcal{M}}_h(P_n))$. Since $\hat{\mathcal{M}}_h(P_n)$ is finite, the Carathéodory theorem in \mathbb{R}^d states that its convex hull is compact.

Note that we can release the assumption $P(\partial H(v, c_{P,h}(v))) = 0$ and $P(\partial\mathcal{B}(x, \delta_{P,h}(x))) = 0$. For this, we proceed like in the proof of Lemma 4.1. According to the Prokhorov lemma and the aforementioned results, we can take a subsequence $Q_n = P_{x_n, h}$ converging to a positive measure Q , and such that $x_n \rightarrow x \in \overline{\mathbb{R}^d}$. We proved in the proof of Lemma 4.1 that Q is necessarily a Borel probability measure such that hQ is a sub-measure of P . It remains to prove that hQ is supported on $\overline{\mathcal{B}(x, \delta_{P,h}(x))}$ and coincides with P on $\mathcal{B}(x, \delta_{P,h}(x))$ if $x \in \mathbb{R}^d$, or that hQ is supported on $\overline{H(v, c_{P,h}(v))}$ and coincides with P on $H(v, c_{P,h}(v))$ if $x \in \overline{\mathbb{R}^d} \setminus \mathbb{R}^d$.

For instance, treat the case $x \in \mathbb{R}^d$. For such an x , we noted that for all $y \notin \partial\mathcal{B}(x, \delta_{P,h}(x))$, $\mathbb{1}_{\mathcal{B}(x_n, \delta_{P,h}(x_n))}(y) \rightarrow \mathbb{1}_{\mathcal{B}(x, \delta_{P,h}(x))}(y)$. Convexity of balls implies that for all $\epsilon > 0$, for n large enough, $\overline{\mathcal{B}(x_n, \delta_{P,h}(x_n))} \subset \overline{\mathcal{B}(x, \delta_{P,h}(x) + \epsilon)}$ and $\overline{\mathcal{B}(x_n, \delta_{P,h}(x_n))} \supset \overline{\mathcal{B}(x, \delta_{P,h}(x) - \epsilon)}$. Thus,

$$\begin{aligned} Q(\overline{\mathcal{B}(x, \delta_{P,h}(x))}) &= \lim_{\epsilon \rightarrow 0} Q(\overline{\mathcal{B}(x, \delta_{P,h}(x) + \epsilon)}) \\ &\geq \lim_{\epsilon \rightarrow 0} \limsup_{n \rightarrow \infty} Q_n(\overline{\mathcal{B}(x, \delta_{P,h}(x) + \epsilon)}) \\ &\geq 1 \end{aligned}$$

since Q_n is supported on $\overline{\mathcal{B}(x_n, \delta_{P,h}(x_n))}$. We also applied the Pormanteau Lemma. Thus, Q is supported on $\overline{\mathcal{B}(x, \delta_{P,h}(x))}$.

Moreover, for all $\eta > 0$,

$$\begin{aligned} Q(\mathcal{B}(x, \delta_{P,h}(x) - \eta)) &= \lim_{\epsilon \rightarrow 0} Q(\overline{\mathcal{B}(x, \delta_{P,h}(x) - \epsilon - \eta)}) \\ &\geq \lim_{\epsilon \rightarrow 0} \limsup_{n \rightarrow \infty} Q_n(\overline{\mathcal{B}(x, \delta_{P,h}(x) - \epsilon - \eta)}) \\ &= \lim_{\epsilon \rightarrow 0} \frac{1}{h} P(\overline{\mathcal{B}(x, \delta_{P,h}(x) - \epsilon - \eta)}) \\ &= \frac{1}{h} P(\mathcal{B}(x, \delta_{P,h}(x) - \eta)). \end{aligned}$$

Thus, $Q(\mathcal{B}(x, \delta_{P,h}(x))) \geq \frac{1}{h} P(\mathcal{B}(x, \delta_{P,h}(x)))$. And hQ and P coincide on $\mathcal{B}(x, \delta_{P,h}(x))$. Thus, $Q \in \mathcal{P}_{x,h}(P)$. We can apply the same method when $x \in \overline{\mathbb{R}^d} \setminus \mathbb{R}^d$, since the support of P is compact.

To conclude, we proceed like in the proof of Lemma 4.1 and get that $l = m(Q)$.

Proof of Lemma 4.19

Let $P \in \mathcal{P}^K(\mathbb{R}^d)$ for some $K > 0$, that is $\text{Supp}(P) \subset \mathcal{B}(0, K)$. For $n \in \mathbb{N}^*$, set $x_n = \lambda n$. For $n \geq c_{P,h}(v)$, we get:

$$\mathcal{B}(x_n, n - c_{P,h}(v)) \subset \mathbb{H}(v, c_{P,h}(v)).$$

Thus, since $P(\mathbb{H}(v, c_{P,h}(v))) \leq h$,

$$\mathcal{B}(x_n, n - c_{P,h}(v)) \subset \mathcal{B}(x_n, \delta_{P,h}(x_n)). \quad (4.26)$$

Moreover, for all $\epsilon > 0$, since $\text{Supp}(P) \subset \mathcal{B}(0, K)$, we get:

$$\mathbb{H}(v, c_{P,h}(v) - \epsilon) \cap \text{Supp}(P) \subset \mathcal{B}\left(x_n, \sqrt{K^2 - 2(c_{P,h}(v) - \epsilon)n + n^2}\right).$$

Thus, since $P(\mathbb{H}(v, c_{P,h}(v) - \epsilon) \cap \text{Supp}(P)) > h$,

$$\mathcal{B}(x_n, \delta_{P,h}(x_n)) \subset \mathcal{B}\left(x_n, \sqrt{K^2 - c_{P,h}^2(v) + (n - c_{P,h}(v))^2}\right). \quad (4.27)$$

Finally, for all $y \in \mathbb{R}^d$, if $\langle y, v \rangle = c_{P,h}(v) - \epsilon$ for some $\epsilon > 0$, then $\|y - x_n\|^2 = \|y\|^2 + n^2 - 2n(c_{P,h}(v) - \epsilon)$, which is superior to $K^2 + (n - c_{P,h}(v))^2 - c_{P,h}^2(v)$ for n large enough. Thus, for all n large enough, $y \notin \mathcal{B}(x_n, \delta_{P,h}(x_n))$. If $\langle y, v \rangle = c_{P,h}(v) + \epsilon$ for some $\epsilon > 0$, then $\|y - x_n\|^2 = \|y\|^2 + (n - c_{P,h}(v))^2 - c_{P,h}^2(v) - 2n\epsilon$ which is inferior to $(n - c_{P,h}(v))^2$ for n large enough. Thus, for all n large enough, $y \in \mathcal{B}(x_n, \delta_{P,h}(x_n))$, which concludes the first part of the Lemma.

Let $(x_n)_{n \geq 0}$ be a sequence in \mathbb{R}^d such that $\lim_{n \rightarrow +\infty} d_{\mathbb{R}^d}(x_n, v_\infty) = 0$, that is such that $\lim_{n \rightarrow +\infty} \|x_n\| = +\infty$ and $\lim_{n \rightarrow +\infty} \frac{x_n}{\|x_n\|} = v$. Then,

$$\|x_n\| - K \leq \delta_{P,h}(x_n) \leq \|x_n\| + K.$$

Let $y \in \mathbb{R}^d$. Then,

$$\|y - x_n\|^2 - \delta_{P,h}(x_n)^2 = \|y\|^2 - 2\langle x_n, y \rangle + O(\|x_n\|) = \|x_n\| \left(\frac{\|y\|^2}{\|x_n\|} - 2 \left\langle \frac{x_n}{\|x_n\|}, y \right\rangle + O(1) \right).$$

The notation $y_n = O(\|x_n\|)$ means that $\left(\frac{y_n}{\|x_n\|}\right)_{n \in \mathbb{N}}$ is bounded. Thus, up to a subsequence,

$$\lim_{n \rightarrow +\infty} \frac{\|y - x_n\|^2 - \delta_{P,h}(x_n)^2}{\|x_n\|} = 2c - 2\langle v, y \rangle,$$

for some $c \in \mathbb{R}$. We deduce that, for all $y \in \mathbb{R}^d \setminus \partial\mathbb{H}(v, c)$,

$$\mathbb{1}_{\mathcal{B}(x_n, \delta_{P,h}(x_n))}(y) \rightarrow \mathbb{1}_{\mathbb{H}(v, c)}(y).$$

In particular, the dominated convergence theorem yields $P(\mathbb{H}(v, c)) \leq h$ and $P(\overline{\mathbb{H}}(v, c)) \geq h$. Therefore, for P -almost y , $\mathbb{1}_{\mathbb{H}(v, c)}(y) = \mathbb{1}_{\mathbb{H}(v, c_{P,h}(v))}(y)$, the result then holds for $c = c_{P,h}(v)$.

Proof of Lemma 4.21

Recall that an *extreme point* x of a convex set S is a point $x \in S$ such that $S \setminus \{x\}$ is still convex. We will prove that the set of extreme points of $\text{Conv}(\overline{\mathcal{M}}_h(P))$ is a subset of $\{m(P_{x,h}) \mid x \in \overline{\mathbb{R}}^d \setminus \mathbb{R}^d\}$. More precisely, we will prove that:

$$\text{Conv}(\overline{\mathcal{M}}_h(P)) = \bigcap_{v \in S(0,1)} H^c(v, \langle m(P_{v_\infty, h}), v \rangle).$$

Set z , an extreme point of $\text{Conv}(\overline{\mathcal{M}}_h(P))$. By definition of the convex hull, $\text{Conv}(\overline{\mathcal{M}}_h(P))$ is the smallest convex set (for the relation of inclusion) containing $\overline{\mathcal{M}}_h(P)$. Thus, necessarily, $z \in \overline{\mathcal{M}}_h(P)$ and $z = m(P_{x,h})$ for some $x \in \overline{\mathbb{R}}^d$. Since z is extreme, $z \in \partial \text{Conv}(\overline{\mathcal{M}}_h(P))$ and according to the Hahn-Banach theorem, there is some vector v and some constant C_x such that $\langle m(P_{x,h}), v \rangle = C_x$ and such that for all $y \in \overline{\mathbb{R}}^d$, $\langle m(P_{y,h}), v \rangle \leq C_x$. We aim at proving that $z = m(P_{v_\infty, h})$ for some $P_{v_\infty, h} \in \mathcal{P}_{v_\infty, h}(P)$.

We can write $P_{x,h}$ as $P_1 + P_2$ with, for all Borel set $B \in \Sigma$, $P_1(B) = P_{x,h}(B \cap H(v, c_{P,h}(v)))$ and $P_2(B) = P_{x,h}(B \cap H^c(v, c_{P,h}(v)))$. Note that P_1 is also a sub-measure of $P_{v_\infty, h}$. Set $P'_2 = P_{v_\infty, h} - P_1$. Then, we get

$$\begin{aligned} C_x &= P_{x,h} \langle u, v \rangle \\ &= P_1 \langle u, v \rangle + P_2 \langle u, v \rangle \\ &= P_{v_\infty, h} \langle u, v \rangle - P'_2 \langle u, v \rangle + P_2 \langle u, v \rangle \\ &\leq C_x - P'_2 \langle u, v \rangle + P_2 \langle u, v \rangle \\ &\leq C_x - c_{P,h}(v) P'_2(\mathbb{R}^d) + c_{P,h}(v) P_2(\mathbb{R}^d) \\ &= C_x \end{aligned}$$

since P_2 is supported on $H^c(v, c_{P,h}(v))$, P'_2 is supported on $\overline{H}(v, c_{P,h}(v))$ and $P_2(\mathbb{R}^d) = P'_2(\mathbb{R}^d)$.

Thus the inequalities are equalities and we get that for P_2 -almost all y , $\langle y, v \rangle = c_{P,h}(v)$ and for P'_2 -almost all y , $\langle y, v \rangle = c_{P,h}(v)$. Thus, $P_{x,h}$ belongs to $\mathcal{P}_{v_\infty, h}(P)$. Moreover, the equalities yield that $C_x = \langle m(P_{x,h}), v \rangle = \langle m(P_{v_\infty, h}), v \rangle$.

Since the set of extreme points is included in $\{m(P_{v_\infty, h}) \mid v \in S(0, 1)\}$, the Krein-Milman theorem yields

$$\text{Conv}(\overline{\mathcal{M}}_h(P)) = \text{Conv}(\{m(P_{v_\infty, h}) \mid v \in S(0, 1)\}).$$

For $v \in S(0, 1)$, set $\overline{C}(P_{v_\infty, h}) = \langle m(P_{v_\infty, h}), v \rangle$, that is, $\overline{C}(P_{v_\infty, h}) = P_{v_\infty, h} \langle u, v \rangle$. For all $x \in \overline{\mathbb{R}}^d$, set P_1 , P_2 and P'_2 as above.

Then, we have

$$\begin{aligned} P_{x,h} \langle u, v \rangle &= P_1 \langle u, v \rangle + P_2 \langle u, v \rangle \\ &= P_{v_\infty, h} \langle u, v \rangle - P'_2 \langle u, v \rangle + P_2 \langle u, v \rangle \\ &= \overline{C}(P_{v_\infty, h}) - \langle m(P'_2), v \rangle P'_2(\mathbb{R}^d) + \langle m(P_2), v \rangle P_2(\mathbb{R}^d) \\ &\leq \overline{C}(P_{v_\infty, h}) - c_{P,h}(v) P'_2(\mathbb{R}^d) + c_{P,h}(v) P_2(\mathbb{R}^d) \\ &= \overline{C}(P_{v_\infty, h}). \end{aligned}$$

Thus, for all $x \in \overline{\mathbb{R}}^d$, $\langle m(P_{x,h}), v \rangle \leq \overline{C}(P_{v_\infty, h}) = \langle m(P_{v_\infty, h}), v \rangle$. In particular, the hyperplane $\partial H(v, \langle m(P_{v_\infty, h}), v \rangle)$ separates $m(P_{v_\infty, h})$ from $\text{Conv}(\overline{\mathcal{M}}_h(P))$.

We proved that for all $y \in \mathbb{R}^d$, for all $v \in S(0, 1)$,

$$\langle m(P_{-v_\infty, h}), v \rangle \leq \langle m(P_{y, h}), v \rangle \leq \langle m(P_{v_\infty, h}), v \rangle.$$

Therefore, the convex set $\text{Conv}(\overline{\mathcal{M}}_h(P))$ is included in $\bigcap_{v \in S(0, 1)} H^c(v, \langle m(P_{v_\infty, h}), v \rangle)$.

Reciprocally, the Hahn-Banach separation theorem applied to the set $\text{Conv}(\overline{\mathcal{M}}_h(P))$ which is compact according to Lemma 4.20 yields that for all $\theta \notin \text{Conv}(\overline{\mathcal{M}}_h(P))$, there is some vector v such that for all $\theta' \in \text{Conv}(\overline{\mathcal{M}}_h(P))$, $\langle \theta', v \rangle \leq C < \langle \theta, v \rangle$. In particular, we get that $\langle \theta, v \rangle > \langle m(P_{v_\infty, h}), v \rangle$, meaning that θ does not belong to $H^c(v, \langle m(P_{v_\infty, h}), v \rangle)$.

Proof of Lemma 4.22

Let $h \in (0, 1)$ and $P \in \mathcal{P}^K(\mathbb{R}^d)$. Then, according to Lemma 4.21, we have:

$$\text{Conv}(\overline{\mathcal{M}}_h(P)) = \bigcap_{v \in S(0, 1)} H^c(v, \langle m(P_{v_\infty, h}), v \rangle).$$

The same holds for h' .

Thus, in order to prove that $\text{Conv}(\overline{\mathcal{M}}_{h'}(P_n)) \subset \text{Conv}(\overline{\mathcal{M}}_h(P_n))$, it is sufficient to prove that

$$H^c(v, \langle m(P_{v_\infty, h'}), v \rangle) \subset H^c(v, \langle m(P_{v_\infty, h}), v \rangle).$$

Thus, it is sufficient to prove that

$$\langle m(P_{v_\infty, h'}), v \rangle \leq \langle m(P_{v_\infty, h}), v \rangle.$$

Set P_0 the sub-measure of P supported on $\overline{H}(v, c_{P, h'}(v)) \setminus H(v, c_{P, h}(v))$ such that $h'P_{v_\infty, h'} = hP_{v_\infty, h} + (h' - h)P_0$. Then, we have:

$$\langle m(P_{v_\infty, h'}), v \rangle = \frac{h}{h'} \langle m(P_{v_\infty, h}), v \rangle + \frac{h' - h}{h'} \langle m(P_0), v \rangle.$$

The results follows from the inequality $\langle m(P_0), v \rangle \leq c_{P, h}(v) \leq \langle m(P_{v_\infty, h}), v \rangle$.

Proof of Lemma 4.23

Let

$$\hat{\mathcal{M}}_h(P_n) = \left\{ \bar{x} = \frac{1}{q} \sum_{p \in \text{NN}_{q, \mathbb{X}_n}(x)} p \mid x \in \overline{\mathbb{R}}^d, \text{NN}_{q, \mathbb{X}_n}(x) \in \mathcal{N}_{q, \mathbb{X}_n}(x) \right\},$$

with $\mathcal{N}_{q, \mathbb{X}_n}(x)$ the collection of all sets of q -nearest neighbors associated to x . Note that different \bar{x} may be associated to the same x , and also note that $\hat{\mathcal{M}}_h(P_n) \subset \overline{\mathcal{M}}_h(P_n)$. Moreover, $\text{Conv}(\overline{\mathcal{M}}_h(P_n)) = \text{Conv}(\hat{\mathcal{M}}_h(P_n))$ since any $m(P_{n, x, h})$, for $P_{n, x, h} \in \mathcal{P}_{x, h}(P_n)$, can be expressed as a convex combination of the \bar{x} 's.

Then, \mathbb{R}^d breaks down into a finite number of weighted Voronoï cells $\mathcal{C}_{P_n, h}(\bar{x}) = \{z \in \mathbb{R}^d \mid \|z - \bar{x}\|^2 + \hat{\omega}^2(\bar{x}) \leq \|z - \bar{y}\|^2 + \hat{\omega}^2(\bar{y}), \forall \bar{y} \in \hat{\mathcal{M}}_h(P_n)\}$, with $\hat{\omega}^2(\bar{x}) = \frac{1}{q} \sum_{p \in \text{NN}_{q, \mathbb{X}_n}(x)} \|p - \bar{x}\|^2$ the weight associated to any point $\bar{x} = \frac{1}{q} \sum_{p \in \text{NN}_{q, \mathbb{X}_n}(x)} p$ in $\hat{\mathcal{M}}_h(P_n)$. According to [BCY17, Theorem 4.3], the weighted Delaunay triangulation partitions the convex hull of any finite set of weighted points \mathbb{X} in general position by d -dimensional simplices with vertices in \mathbb{X} , provided that the associated weighted Voronoï cells of all the points in \mathbb{X} are non empty. By duality, (also see [BCY17, Lemma 4.5]) these vertices are associated to weighted Voronoï cells that have non-empty common intersection. Thus, any $\theta \in \text{Conv}(\overline{\mathcal{M}}_h(P_n))$ satisfies $\theta = \sum_{i=0}^d \lambda_i \bar{x}^i$ for some \bar{x}^i 's in

$\hat{\mathcal{M}}_h(P_n)$ and some non negative λ_i 's such that $\sum_{i=0}^d \lambda_i = 1$. Also, there exists some x^* in the intersection of the $d + 1$ weighted Voronoï cells, $(\mathcal{C}_{P_n, h}(\bar{x}^i))_{i \in \llbracket 0, d \rrbracket}$. Set $P_{n x^*, h} := \sum_{i=0}^d \lambda_i P_i$, with $P_i = \frac{1}{q} \sum_{p \in \text{NN}_{q, \bar{x}_n}^i(x^*)} \delta_{\{p\}}$ when $\bar{x}^i = \frac{1}{q} \sum_{p \in \text{NN}_{q, \bar{x}_n}^i(x^*)} p$. Then, $P_{n x^*, h}$ is a probability measure such that $h P_{n x^*, h}$ ($h = \frac{q}{n}$) is a sub-measure of P_n , coincides with P_n on $\mathcal{B}(x, \delta_{P_n, h}(x))$ and is supported on $\overline{\mathcal{B}}(x, \delta_{P_n, h}(x))$. Thus it belongs to $\mathcal{P}_{x^*, h}(P_n)$. Moreover, its mean $m(P_{n x^*, h}) = \theta$. Thus, $\theta \in \overline{\mathcal{M}}_h(P_n)$.

Proof of Lemma 4.24

The proof of Lemma 4.24 is based on the concentration argument given in Lemma 3.58, that allows to connect empirical sub-measures with sub-measures for P .

A significant part of the proof of Lemma 4.24 is based on the characterization of $\text{Conv}(\overline{\mathcal{M}}_h(P))$ through $\omega_{P, h}^2$ stated by Lemma 4.33.

Lemma 4.40. *Let C denote a convex set, $\theta \in \mathbb{R}^d$, and $\Delta = d(\theta, C)$. There exists $v \in \mathbb{R}^d$ with $\|v\| = 1$ such that, for all τ in C ,*

$$\langle v, \theta - \tau \rangle \geq \Delta.$$

Proof of Lemma 4.40. Denote by π the projection onto C , and $t = \pi(\theta)$. Then, let $x = \frac{\theta - t}{\Delta}$. We may write

$$\begin{aligned} \langle x, \theta - \tau \rangle &= \langle x, \theta - t \rangle + \langle x, t - \tau \rangle \\ &= \Delta + \frac{1}{\Delta} \langle \theta - t, t - \tau \rangle. \end{aligned}$$

Since, for all τ in C , $\langle \theta - t, \tau - t \rangle \leq 0$, the result follows. □

We are now in position to prove Lemma 4.24.

Proof of Lemma 4.24. Let P in $\mathcal{P}^K(\mathbb{R}^d)$ which puts mass neither on balls nor on hyperplanes. For some $h \in (0, 1)$, let $\theta \in \text{Conv}(\overline{\mathcal{M}}_h(P))$. For p large enough (for instance $p = 10$), a union bound ensures that the inequalities of Lemma 3.58 are satisfied for all $n \geq 2$ with positive probability. According to the assumption of the Lemma, there are $X_1, X_2, \dots, X_n, \dots$ for which the points in $\hat{\mathcal{M}}_h(P_n)$ are in general position for all $h = \frac{q}{n}$ for $n \in \mathbb{N}^*$ and $q \in \llbracket 1, n - 1 \rrbracket$, and such that the inequalities of Lemma 3.58 are satisfied for all $n \geq 2$. In particular, for $n \in \mathbb{N}^*$, for such distributions P_n and (y, r) such that $P(\mathcal{B}(y, r)) = h$, we have

$$\begin{aligned} \left\| \frac{P_n u \mathbb{1}_{\mathcal{B}(y, r)}(u)}{P(\mathcal{B}(y, r))} - \frac{P_n u \mathbb{1}_{\mathcal{B}(y, r)}(u)}{P_n(\mathcal{B}(y, r))} \right\| &\leq \frac{K \alpha_n}{h}, \\ \left\| \frac{P u \mathbb{1}_{\mathcal{B}(y, r)}(u)}{P(\mathcal{B}(y, r))} - \frac{P_n u \mathbb{1}_{\mathcal{B}(y, r)}(u)}{P(\mathcal{B}(y, r))} \right\| &\leq \frac{K \alpha_n}{h}, \\ |(P_n - P)\mathcal{B}(y, r)| &\leq \alpha_n, \end{aligned}$$

for $\alpha_n \rightarrow 0$. Note that the same holds for means on half-spaces. Now let $x \in \mathbb{R}^d$,

$$\begin{aligned} d_{P, h}^2(x) - \|x - \theta\|^2 &= P_{x, h} \|x - u\|^2 - \|x - \theta\|^2 \\ &= \inf_{y \in \overline{\mathbb{R}}^d} P_{y, h} \|x - u\|^2 - \|x - \theta\|^2 \\ &\geq \inf_{y \in \overline{\mathbb{R}}^d} \|m(P_{y, h})\|^2 + v(P_{y, h}) - \|\theta\|^2 + \inf_{y \in \overline{\mathbb{R}}^d} 2\langle x, \theta - m(P_{y, h}) \rangle \\ &\geq -\|\theta\|^2 + \inf_{y \in \overline{\mathbb{R}}^d} 2\langle x, \theta - m(P_{y, h}) \rangle. \end{aligned}$$

Thus, we may write

$$\begin{aligned} \inf_{y \in \mathbb{R}^d} 2 \langle x, \theta - m(P_{y,h}) \rangle &\geq \min \left[\inf_{y,r | P_n(\mathcal{B}(y,r)) \in [h-\alpha_n, h+\alpha_n]} 2 \left\langle x, \theta - \frac{P_n u \mathbb{1}_{\mathcal{B}(y,r)}(u)}{P_n(\mathcal{B}(y,r))} \right\rangle, \right. \\ &\quad \left. \inf_{v,c | P_n(H(v,c)) \in [h-\alpha_n, h+\alpha_n]} 2 \left\langle x, \theta - \frac{P_n u \mathbb{1}_{H(v,c)}(u)}{P_n(H(v,c))} \right\rangle \right] - \frac{4K\alpha_n \|x\|}{h}, \\ &= \inf_{\tau \in \bigcup_{s \in [h-\alpha_n, h+\alpha_n]} \overline{\mathcal{M}}_s(P_n)} 2 \langle x, \theta - \tau \rangle - \frac{4K\alpha_n \|x\|}{h}. \end{aligned}$$

Now, if $d \left(\theta, \text{Conv} \left(\bigcup_{s \in [h-\alpha_n, h+\alpha_n]} \overline{\mathcal{M}}_s(P_n) \right) \right) = \Delta > \frac{2K\alpha_n}{h}$, then according to Lemma 4.40, we can choose x in \mathbb{R}^d such that, for all $\tau \in \text{Conv} \left(\bigcup_{s \in [h-\alpha_n, h+\alpha_n]} \overline{\mathcal{M}}_s(P_n) \right)$,

$$\left\langle \frac{x}{\|x\|}, \theta - \tau \right\rangle - \frac{2K\alpha_n}{h} > 0.$$

In this case, we immediately get $\omega_{P,h}^2(\theta) = \sup_{x \in \mathbb{R}^d} d_{P,h}^2(x) - \|x - \theta\|^2 = +\infty$. According to Lemma 4.33, this contradicts $\theta \in \text{Conv}(\overline{\mathcal{M}}_h(P))$.

Set $h_n = \frac{q_n}{n}$ for $q_n \in \llbracket 1, n-1 \rrbracket$ such that $h - \alpha_n \geq h_n \geq h - \alpha_n - \frac{1}{n}$. Note that for n large enough, $h - \alpha_n - \frac{1}{n} > 0$, thus h_n is well defined. Then, according to Lemma 4.22 and 4.23,

$$\text{Conv} \left(\bigcup_{s \in [h-\alpha_n, h+\alpha_n]} \overline{\mathcal{M}}_s(P_n) \right) \subset \text{Conv} \left(\bigcup_{s \in [h_n, 1]} \overline{\mathcal{M}}_s(P_n) \right) = \overline{\mathcal{M}}_{h_n}(P_n).$$

Thus, we can build a sequence $(y_n)_{n \geq N}$ for some $N \in \mathbb{N}$ such that $y_n \in \overline{\mathcal{M}}_{h_n}(P_n)$ and $\|\theta - y_n\| \leq 2\frac{K\alpha_n}{h}$. Hence the result of Lemma 4.24. \square

Proof of Lemma 4.33

According to Proposition 4.18, for all $x \in \mathbb{R}^d$, $d_{P,h}^2(x) - \|x - \theta\|^2$ may be written as

$$\inf_{y \in \mathbb{R}^d} \inf_{P_{y,h} \in \mathcal{P}_{y,h}(P)} \left\{ \|m(P_{y,h})\|^2 + v(P_{y,h}) - \|\theta\|^2 + 2 \langle x, \theta - m(P_{y,h}) \rangle \right\},$$

which is lower-bounded by

$$\inf_{y \in \mathbb{R}^d} \inf_{P_{y,h} \in \mathcal{P}_{y,h}(P)} \left\{ \|m(P_{y,h})\|^2 + v(P_{y,h}) \right\} - \|\theta\|^2 + \inf_{\tau \in \overline{\mathcal{M}}_h(P)} \left\{ 2 \langle x, \theta - \tau \rangle \right\}.$$

Assume $\theta \notin \text{Conv}(\overline{\mathcal{M}}_h(P))$. According to Lemma 4.20, $\text{Conv}(\overline{\mathcal{M}}_h(P))$ is a convex and compact subset of \mathbb{R}^d . The Hahn-Banach separation theorem thus provides some vector $v \in \mathbb{R}^d$ and $C < 0$ such that $\forall \tau \in \overline{\mathcal{M}}_h(P)$, $\langle \theta - \tau, v \rangle < C$. Setting $x_n = -nv$ for $n \in \mathbb{N}^*$ yields $\lim_{n \rightarrow +\infty} \inf_{\tau \in \overline{\mathcal{M}}_h(P)} \langle x_n, \theta - \tau \rangle = +\infty$. Thus, $\sup_{x \in \mathbb{R}^d} d_{P,h}^2(x) - \|x - \theta\|^2 = +\infty$.

Now, let $\theta \in \text{Conv}(\overline{\mathcal{M}}_h(P))$, we can write $\theta = \sum_{i=0}^d \lambda_i m(P_i)$ for $P_i = P_{x_i, h}$ with the x_i 's in $\overline{\mathbb{R}}^d$. We have:

$$\begin{aligned} \sup_{x \in \overline{\mathbb{R}}^d} d_{P, h}^2(x) - \|x - \theta\|^2 &= \sup_{x \in \overline{\mathbb{R}}^d} \sum_{i=0}^d \lambda_i (d_{P, h}^2(x) - \|x - \theta\|^2) \\ &\leq \sup_{x \in \overline{\mathbb{R}}^d} \sum_{i=0}^d \lambda_i (\|x - m(P_i)\|^2 + v(P_i) - \|x - \theta\|^2) \\ &= \sup_{x \in \overline{\mathbb{R}}^d} \sum_{i=0}^d \lambda_i (v(P_i) + 2\langle x, \theta - m(P_i) \rangle + \|m(P_i)\|^2 - \|\theta\|^2) \\ &= \sum_{i=0}^d \lambda_i (v(P_i) + \|m(P_i)\|^2 - \|\theta\|^2), \end{aligned}$$

according to Proposition 4.18. Thus, we get that

$$\omega_{P, h}^2(\theta) + \|\theta\|^2 \leq \sum_{i=0}^d \lambda_i (v(P_i) + \|m(P_i)\|^2) \leq \sup_{x \in \overline{\mathbb{R}}^d} \{v(P_{x, h}) + \|m(P_{x, h})\|^2\}. \quad (4.28)$$

Lemma 4.36 yields $\omega_{P, h}^2(\theta) < +\infty$.

Bibliography

- [Aam17] Eddie Aamari. *Vitesses de convergence en inférence géométrique*. PhD thesis, Université Paris-Saclay, 2017.
- [AB06] Charalambos D. Aliprantis and Kim C. Border. *Infinite Dimensional Analysis: a Hitchhiker's Guide*. Springer, Berlin; London, 2006.
- [AC17] Catherine Aaron and Alejandro Cholaquidis. On boundary detection. working paper or preprint, July 2017.
- [AE75] David A. Edwards. The structure of superspace. In Nick M. Stavrakas and Keith R. Allen, editors, *Studies in Topology*, pages 121–133. Academic Press, 12 1975.
- [AG80] Aloisio Araujo and Evarist Giné. *The Central Limit Theorem for Real and Banach Valued Random Variables*. John Wiley & Sons Inc, 1980.
- [Arl07] Sylvain Arlot. *Rééchantillonnage et Sélection de Modèles*. PhD thesis, Université Paris Sud, 2007.
- [AT11] Masafumi Akahira and Kei Takeuchi. *Asymptotic Efficiency of Statistical Estimators: Concepts and Higher Order Asymptotic Efficiency: Concepts and Higher Order Asymptotic Efficiency*. Lecture Notes in Statistics. Springer New York, 2011.
- [AV07] David Arthur and Sergei Vassilvitskii. k-means++: the advantages of careful seeding. In *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1027–1035. ACM, New York, 2007.
- [BBL05] Stéphane Boucheron, Olivier Bousquet, and Gábor Lugosi. Theory of classification: A survey of some recent advances. *ESAIM: Probability and Statistics*, 9:323, 2005.
- [BC93] Tadeusz Bednarski and Brenton R. Clarke. Trimmed likelihood estimation of location and scale of the normal distribution. *Australian Journal of Statistics*, 35(2):141–153, 1993.

- [BCOS15] Mickaël Buchet, Frédéric Chazal, Steve Y. Oudot, and Donald R. Sheehy. Efficient and robust persistent homology for measures. In *Proceedings of the Twenty-sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '15*, pages 168–180, Philadelphia, PA, USA, 2015. Society for Industrial and Applied Mathematics.
- [BCRR97] Halima Bensmail, Gilles Celeux, E. Adrian Raftery, and Christian P. Robert. Inference in model-based cluster analysis. *Statistics and Computing*, 7:1–10, 1997.
- [BCY17] Jean-Daniel Boissonnat, Frédéric Chazal, and Mariette Yvinec. *Geometric and Topological Inference*. Cambridge University Press, October 2017. To appear in the red series Cambridge Texts in Applied Mathematics.
- [BDL08] Gérard Biau, Luc Devroye, and Gábor Lugosi. On the performance of clustering in hilbert spaces. *IEEE Trans. Information Theory*, 54(2):781–790, 2008.
- [BFL18] Claire Brécheteau, Aurélie Fischer, and Clément Levrard. Robust bregman clustering. Submitted to NIPS 2018, 2018.
- [BGW05] Arindam Banerjee, Xin Guo, and Hui Wang. On the optimality of conditional expectation as a bregman predictor. *IEEE Transactions on Information Theory*, 51(7):2664 – 2669, 2005.
- [Bil99] Patrick Billingsley. *Convergence of Probability Measures*. Wiley-Interscience, 1999.
- [BK04] Mireille Boutin and Gregor Kemper. On reconstructing n-point configurations from the distribution of distances or areas. *Advances in Applied Mathematics*, 32:709–735, 2004.
- [BL14] Sergey Bobkov and Michel Ledoux. One-dimensional empirical measures, order statistics and kantorovich transport distances. *the Memoirs of the Amer. Math. Soc.*, 2014. to be published.
- [BL16] Blanche Buet and Gian Paolo Leonardi. Recovering measures from approximate values on balls. *Annales Academiae Scientiarum Fennicae Mathematica*, 41, 2016.
- [BL17] Claire Brécheteau and Clément Levrard. The k-pdtm: A coresets for robust geometric inference. Unpublished, 2017.
- [BLM13] Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration Inequalities: A Nonasymptotic Theory of Independence*. OUP Oxford, 2013.
- [Blo77] Gary S Bloom. A counterexample to a theorem of s.piccard. *Journal of Combinatorial Theory, Series A*, 22:378–379, 1977.
- [BMDG05] Arindam Banerjee, Srujana Merugu, Inderjit S. Dhillon, and Joydeep Ghosh. Clustering with bregman divergences. *J. Mach. Learn. Res.*, 6:1705–1749, December 2005.
- [BNN10] Jean-Daniel Boissonnat, Frank Nielsen, and Richard Nock. Bregman voronoi diagrams. *Discrete & Computational Geometry*, 44(2):281–307, 2010.
- [Bor98] Émile Borel. *Leçons sur la théorie des fonctions*. Collection de monographies sur la théorie des fonctions. Gauthier-Villars et fils, 1898.

- [Bor30] Émile Borel. *Traité du calcul des probabilités et de ses applications*, volume I of *Traité du calcul des probabilités et de ses applications*. Gautier-Villars et cie, 1930.
- [Bre67] Lev M. Bregman. The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming. *USSR Computational Mathematics and Mathematical Physics*, 7(3):200 – 217, 1967.
- [Bré16] Claire Brécheteau. The dtm-signature for a geometric comparison of metric-measure spaces from samples. Unpublished, 2016.
- [Bré18] Claire Brécheteau. A statistical test of isomorphism between metric-measure spaces using the distance-to-a-measure signatures. Submitted to *Electronic Journal of Statistics*, 2018.
- [Buc14] Mickaël Buchet. *Topological inference from measures*. PhD thesis, Université Paris Sud - Paris XI, 2014.
- [CAGM97] Juan. A. Cuesta-Albertos, Alfonso Gordaliza, and Carlos Matrán. Trimmed k -means: an attempt to robustify quantizers. *Ann. Statist.*, 25(2):553–576, 04 1997.
- [CAMMI08] Juan. A. Cuesta-Albertos, Carlos Matrán, and Agustín Mayo-Isacar. Robust estimation in the normal mixture model based on robust clustering. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 70(4):779–802, 2008.
- [CCSG⁺09] Frédéric Chazal, David Cohen-Steiner, Leonidas J. Guibas, Facundo Mémoli, and Steve Oudot. Gromov-Hausdorff stable signatures for shapes using persistence. *Computer Graphics Forum (proc. SGP 2009)*, pages 1393–1403, 2009.
- [CCSM11] Frédéric Chazal, David Cohen-Steiner, and Quentin Mérigot. Geometric inference for probability measures. *Foundations of Computational Mathematics archive*, 11:733 – 751, 2011.
- [CdSGO16] Frédéric Chazal, Vin de Silva, Marc Glisse, and Steve Y. Oudot. *The Structure and Stability of Persistence Modules*. Springer Briefs in Mathematics. Springer, 2016.
- [CDSO14] Frédéric Chazal, Vin De Silva, and Steve Oudot. Persistence stability for geometric complexes. *Geometriae Dedicata*, 173(1):193–214, 2014.
- [CFL⁺13] Frédéric Chazal, Brittany Terese Fasy, Fabrizio Lecci, Alessandro Rinaldo, Aarti Singh, and Larry Wasserman. On the bootstrap for persistence diagrams and landscapes. *Modeling and Analysis of Information Systems*, 20:96–105, 2013.
- [CG95] Gilles Celeux and Gérard Govaert. Gaussian parsimonious clustering models. *Pattern Recognition*, 28(5):781 – 793, 1995.
- [CGLM15] Frédéric Chazal, Marc Glisse, Catherine Labruère, and Bertrand Michel. Convergence rates for persistence diagram estimation in topological data analysis. *J. Mach. Learn. Res.*, 16(1):3603–3635, January 2015.
- [CGOS13] Frédéric Chazal, Leonidas J. Guibas, Steve Y. Oudot, and Primoz Skraba. Persistence-based clustering in riemannian manifolds. *J. ACM*, 60(6):41:1–41:38, November 2013.
- [CL05] Frédéric Chazal and André Lieutier. The λ -medial axis. *Graphical Models*, 67:304 – 331, 2005.

- [CM74] N. A. Campbell and R. J. Mahon. A multivariate study of variation in two species of rock crab of genus *emphLeptograpsus*. *Australian Journal of Zoology*, 22:417–425, 1974.
- [CMM16] Frédéric Chazal, Pascal Massart, and Bertrand Michel. Rates of convergence for robust geometric inference. *Electron. J. Statist.*, 10(2):2243–2286, 2016.
- [Coh80] D.L. Cohn. *Measure Theory*. Birkhäuser, 1980.
- [CSEH07] David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Stability of persistence diagrams. *Discrete & Computational Geometry*, 37(1):103–120, Jan 2007.
- [dAG79] Alejandro de Acosta and Evarist Giné. *Convergence Of Moments And Related Functionals In The Central Limit Theorem In Banach Spaces*. Z. Wahrsch.ver.Geb., 1979.
- [dBGM99] Eustasio del Barrio, Evarist Giné, and Carlos Matrán. Central limit theorems for the Wasserstein distance between the empirical and the true distributions. *The Annals of Probability*, 27(2):1009–1071, 1999.
- [DBLL15] Eustasio Del Barrio, Hélène Lescornel, and Jean-Michel Loubes. A statistical analysis of a deformation model with wasserstein barycenters : estimation procedure and goodness of fit test. unpublished, 2015.
- [DFGEMI18] Francesco Dotto, Alessio Farcomeni, Luis Angel García-Escudero, and Agustín Mayo-Isacar. A reweighting approach to robust clustering. *Statistics and Computing*, 28(2):477–493, 2018.
- [Dob70] R. L. Dobrushin. Prescribing a system of random variables by conditional distributions. *Theory of Probability & Its Applications*, 15(3):458–486, 1970.
- [DR98] Abhijit Dasgupta and Adrian E. Raftery. Detecting features in spatial point processes with clutter via model-based clustering. *Journal of the American Statistical Association*, 93(441):294–302, 1998.
- [Dud79] R.M Dudley. Balls in r_k do not cut all subsets of $k + 2$ points. *Advances in Mathematics*, 31(3):306 – 308, 1979.
- [EC17] Aamari Eddie and Levrard Clément. Stability and minimax optimality of tangential delaunay complexes for manifold reconstruction. *Discrete and Computational Geometry*, 2017.
- [Ede92] Herbert. Edelsbrunner. *Weighted Alpha Shapes*. Number num 1760 in Report. University of Illinois at Urbana-Champaign, Department of Computer Science, 1992.
- [eE53] R. Fortet et E.Mourier. Convergence de la répartition empirique vers la répartition théorique. *Ann. Scient. École Norm. Sup.*, 70:266–285, 1953.
- [Efr79] B. Efron. Bootstrap methods: Another look at the jackknife. *Annals of Statistics*, 7:1–26, 1979.
- [ELZ02] Edelsbrunner, Letscher, and Zomorodian. Topological persistence and simplification. *Discrete & Computational Geometry*, 28(4):511–533, Nov 2002.

- [Fed59] Herbert Federer. *Curvature measures*, volume 93. Trans. Amer. Math. Soc., 1959.
- [Fed96] Herbert Federer. *Geometric Measure Theory*. Springer-Verlag Berlin Heidelberg, 1996.
- [FG15] Nicolas Fournier and Arnaud Guillin. On the rate of convergence in Wasserstein distance of the empirical measure. *Probability Theory and Related Fields*, 162(3-4):707, August 2015.
- [FGEMI12] Heinrich Fritz, Luis A. García-Escudero, and Agustín Mayo-Iscar. tclust: An R package for a trimming approach to cluster analysis, 2012.
- [Fis22] R.A. Fisher. On the mathematical foundations of theoretical statistics. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 222(594-604):309–368, 1922.
- [Fis10] Aurélie Fischer. Quantization and clustering with bregman divergences. *Journal of Multivariate Analysis*, 101(9):2207 – 2221, 2010.
- [FKLM14] Brittany Terese Fasy, Jisu Kim, Fabrizio Lecci, and Clément Maria. Introduction to the R package TDA. *CoRR*, abs/1411.1830, 2014.
- [FLLRB12] Magalie Fromont, Béatrice Laurent, Matthieu Lerasle, and Patricia Reynaud-Bouret. Kernels based tests with non-asymptotic bootstrap approaches for two-sample problems. In *Journal of Machine Learning Research: Workshop and Conference proceedings*, volume 23, pages 23.1–23.22. COLT 2012, 2012.
- [FLP⁺51] K. Florek, J. Łukaszewicz, J. Perkal, Hugo Steinhaus, and S. Zubrzycki. Sur la liaison et la division des points d’un ensemble fini. *Colloquium Mathematicae*, 2(3–4):282–285, 1951.
- [FLR⁺14] Brittany Terese Fasy, Fabrizio Lecci, Alessandro Rinaldo, Larry Wasserman, Sivaraman Balakrishnan, and Aarti Singh. Confidence sets for persistence diagrams. *Ann. Statist.*, 42(6):2301–2339, 12 2014.
- [For65] E. W. Forgy. Cluster analysis of multivariate data: efficiency versus interpretability of classifications. *Biometrics*, 21:768–769, 1965.
- [FR98] C. Fraley and A. E. Raftery. How many clusters? which clustering method? answers via model-based cluster analysis. *The Computer Journal*, 41:578–588, 1998.
- [FR02] Chris Fraley and Adrian E Raftery. Model-based clustering, discriminant analysis, and density estimation. *Journal of the American Statistical Association*, 97(458):611–631, 2002.
- [G.65] Sheals J. G. The application of computer techniques to acarine taxonomy: a preliminary examination with species of the hypoaspis-androlaelaps complex (acarina). *Proceedings of the Linnean Society of London*, 176(1):11–21, 1965.
- [GEGMM10] L. García-Escudero, A. Gordaliza, C. Matrán, and A. Mayo. A review of robust clustering methods. *Advances in Data Analysis and Classification*, 4:89–109, 2010.
- [GEGMMI08] Luis A. García-Escudero, Alfonso Gordaliza, Carlos Matrán, and Agustín Mayo-Iscar. A general trimming approach to robust cluster analysis. *Ann. Statist.*, 36(3):1324–1345, 06 2008.

- [GL00] Siegfried Graf and Harald Luschgy. *Foundations of Quantization for Probability Distributions*. Springer-Verlag, Berlin, Heidelberg, 2000.
- [GLP81] M. Gromov, J. Lafontaine, and P. Pansu. *Structures métriques pour les variétés riemanniennes*. Textes mathématiques. CEDIC/Fernand Nathan, 1981.
- [GLP99] M. Gromov, J. Lafontaine, and P. Pansu. *Metric Structures for Riemannian and Non-Riemannian Spaces*. Progress in Mathematics - Birkhäuser. Birkhäuser, 1999.
- [GMBJR⁺12] Arthur Gretton, Karsten M. Borgwardt, Malte J. Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *Journal of Machine Learning Research*, 13:723–773, 2012.
- [GMM11] Leonidas J. Guibas, Quentin Mérigot, and Dmitriy Morozov. Witnessed k-distance. In *Proceedings of the Twenty-seventh Annual Symposium on Computational Geometry*, SoCG '11, pages 57–64, New York, NY, USA, 2011. ACM.
- [Gor91] Alfonso Gordaliza. Best approximations to random variables based on trimming procedures. *Journal of Approximation Theory*, 64(2):162 – 180, 1991.
- [GR05] María Teresa Gallegos and Gunter Ritter. A robust method for cluster analysis. *Ann. Statist.*, 33(1):347–380, 02 2005.
- [Gro81] Mikhael Gromov. Groups of polynomial growth and expanding maps. *Publications Mathématiques de l'Institut des Hautes Études Scientifiques*, 53(1):53–78, 1981.
- [Hen04] Christian Hennig. Breakdown points for maximum likelihood estimators of location–scale mixtures. *Ann. Statist.*, 32(4):1313–1340, 08 2004.
- [HL97] Ali S. Hadi and Alberto Luceño. Maximum trimmed likelihood estimators: a unified approach, examples, and algorithms. *Computational Statistics & Data Analysis*, 25(3):251 – 272, 1997.
- [Hub64] Peter J. Huber. Robust estimation of a location parameter. *Ann. Math. Statist.*, 35(1):73–101, 03 1964.
- [Čí08] Pavel Čížek. General trimmed estimation: robust approach to nonlinear and limited dependent variable models. *Econometric Theory*, 24(6):1500—1529, 2008.
- [JJN85] P. Jansser, J. Jureckova, and Veraverbeke N. Rate of convergence of one- and two-step m-estimators with applications to maximum likelihood and pitman estimators. *Annals of Statistics*, 13(3):1222–1229, 1985.
- [Joh67] Stephen Johnson. Hierarchical clustering schemes. *Psychometrika*, 32(3):241–254, 1967.
- [Klo12] Benoît Kloekner. Approximation by finitely supported measures. *ESAIM: Control, Optimisation and Calculus of Variations*, 18(2):343—359, 2012.
- [Kol33] A. N. Kolmogorov. *Grundbegriffe der Wahrscheinlichkeitsrechnung*. Springer, Berlin, 1933.
- [Kom88] Hidetoshi Komiya. Elementary proof for sion's minimax theorem. *Kodai Math. J.*, 11(1):5–7, 1988.

- [KR58] L.V. Kantorovich and G.Sh. Rubinshtein. On a space of completely additive functions. *Vestnik Leningrad. Univ.*, 13:52–59, 1958. En russe.
- [Lév37] Paul Lévy. Théorie de l’addition des variables aléatoires. In *Monographies des probabilités*. Paris : Gautier-Villars, 1937.
- [Lev14] Clément Levrard. *Quantification vectorielle en grande dimension : vitesses de convergence et sélection de variables*. PhD thesis, Université Paris-Sud, 2014.
- [Lev15] Clément Levrard. Nonasymptotic bounds for vector quantization in hilbert spaces. *Ann. Statist.*, 43(2):592–619, 04 2015.
- [Lie04] André Lieutier. Any open bounded subset of \mathbb{R}^n has the same homotopy type as its medial axis. *Computer-Aided Design*, 36:1029 – 1046, 2004.
- [Lin02] Tamás Linder. Learning-theoretic methods in vector quantization. In *Principles of nonparametric learning (Udine, 2001)*, *CISM Courses and Lectures*, volume 434, pages 163—210. Springer, Vienna, 2002.
- [Llo82] S. P. Lloyd. Least squares quantization in PCM. *IEEE Transactions on Information Theory*, 28:129–137, 1982.
- [Lug02] Gabor Lugosi. Pattern classification and learning theory. In *Principles of nonparametric learning (Udine, 2001)* *CISM Courses and Lectures*, volume 434, pages 1—56. Springer, Vienna, 2002.
- [Mü97] Alfred Müller. Integral probability metrics and their generating classes of functions. *Advances in Applied Probability*, 29(2):429—443, 1997.
- [Mac67] J. MacQueen. Some methods for classification and analysis of multivariate observations. In *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Statistics*, pages 281–297, Berkeley, Calif., 1967. University of California Press.
- [Mar00] Marianthi Markatou. Mixture models, robustness, and the weighted likelihood methodology. *Biometrics*, 56(2):483–486, 2000.
- [Mas90] Pascal Massart. The tight constant in the Dvoretzky-Kiefer-Wolfowitz inequality. *The Annals of Probability*, 18(3):1269–1283, 1990.
- [Mas07] Pascal Massart. *Concentration Inequalities and Model Selection*. Springer-Verlag Berlin Heidelberg, 2007.
- [Mém11] Facundo Mémoli. Gromov–Wasserstein distances and the metric approach to object matching. *Foundations of Computational Mathematics*, 11(4):417–487, 2011.
- [Mér13] Quentin Mérigot. Lower bounds for k-distance approximation. In *SoCG ’13 - 29th ACM Symposium on Computational Geometry*, pages 435–440, Rio de Janeiro, Brazil, June 2013. 15 pages.
- [MP00] G. J. McLachlan and D. Peel. *Finite mixture models*. Wiley Series in Probability and Statistics, New York, 2000.
- [Mül01] Ch.H. Müller. Trimmed likelihood estimators in generalized linear models. In *Sixth International Conference on Computer Data Analysis and Modeling*, volume 2, pages 142–150, 2001. eds. S. Aivazian, Y. Kharin, H. Rieder, Minsk.

- [MV03] S. Mendelson and R. Vershynin. Entropy and the combinatorial dimension. *Inventiones mathematicae*, 152:37–55, 2003.
- [NBN07] Frank Nielsen, Jean-Daniel Boissonnat, and Richard Nock. Bregman voronoi diagrams: Properties, algorithms and applications. *CoRR*, abs/0709.2196, 2007.
- [NFDN07] N. M. Neykov, Peter Filzmoser, R. Dimova, and P. N. Neytchev. Robust fitting of mixtures using the trimmed likelihood estimator. *Computational Statistics & Data Analysis*, 52:299–308, 2007.
- [NN90] N. Neykov and Plamen Neytchev. A robust alternative of the maximum likelihood estimators. In *COMPSTAT 1990 - Short Communications*, pages 99–100, 01 1990.
- [NSW08] Partha Niyogi, Steven Smale, and Shmuel Weinberger. Finding the homology of submanifolds with high confidence from random samples. *Discrete and Computational Geometry*, 39:419–441, 2008.
- [OFCD02] Robert Osada, Thomas Funkhouser, Bernard Chazelle, and David Dobkin. Shape distributions. *ACM Transactions on Graphics*, 21:807–832, 2002.
- [PM00] D. Peel and G. J. McLachlan. Robust mixture modelling using the t distribution. *Statistics and Computing*, 10(4):339–348, Oct 2000.
- [Pol82a] David Pollard. A central limit theorem for k-means clustering. *The Annals of Probability*, 10(4):919–26, 1982.
- [Pol82b] David Pollard. Quantization and the method of k -means. *IEEE Transactions on Information Theory*, 28(2), 1982.
- [PR94] Dimitris N. Politis and Joseph P. Romano. Large sample confidence regions based on subsamples under minimal assumptions. *The Annals of Statistics*, 22(4):2031–2050, 1994.
- [Pro56] Y. V. Prokhorov. Convergence of Random Processes and Limit Theorems in Probability Theory. *Theory of Probability and its Applications*, 1(2):157–214, 1956.
- [PWZ15] J. M. Phillips, Bei Wang, and Y. Zheng. Geometric inference on kernel density estimates. *CoRR*, abs/1307.7760, 2015.
- [RAW09] Galen R. Shorack and Jon A. Wellner. *Empirical processes with applications to statistics*. Society for Industrial & Applied Mathematics, 2009.
- [RL87] Peter J. Rousseeuw and Annick M. Leroy. *Robust Regression & Outlier Detection*. Wiley, 1987.
- [Roc70] R. Tyrrell Rockafellar. *Convex analysis*, 1970. Princetown Landmarks in Mathematics. Princetown University Press.
- [Rou84] Peter J. Rousseeuw. Least median of squares regression. *Journal of the American Statistical Association*, 79(388):871–880, 1984.
- [Rou85] P. J. Rousseeuw. Multivariate estimation with high breakdown point. *Mathematical statistics and applications*, 8, 1985.

- [RW84] Richard A. Redner and Homer F. Walker. Mixture densities, maximum likelihood and the em algorithm. *SIAM Review*, 26(2):195–239, 1984.
- [SJ02] A. Strehl and Ghosh J. Cluster ensembles – a knowledge reuse framework for combining multiple partitions. *Journal of Machine Learning Research*, 3:583–617, 2002.
- [Ste56] H. Steinhaus. Sur la division des corp materiels en parties. *Bull. Acad. Polon. Sci*, 1:801–804, 1956.
- [Ste00] Matthew Stephens. Bayesian analysis of mixture models with an unknown number of components—an alternative to reversible jump methods. *Ann. Statist.*, 28(1):40–74, 02 2000.
- [TK18] Qingguo Tang and R. J. Karunamuni. Robust variable selection for finite mixture regression models. *Annals of the Institute of Statistical Mathematics*, 70(3):489–521, Jun 2018.
- [Vap82] Vladimir Vapnik. *Estimation of Dependences Based on Empirical Data: Springer Series in Statistics (Springer Series in Statistics)*. Springer-Verlag, Berlin, Heidelberg, 1982.
- [Var58a] V. S. Varadarajan. On a problem in measure-spaces. *Ann. Math. Statist.*, 29(4):1275–1278, 12 1958.
- [Var58b] V. S. Varadarajan. On the convergence of sample probability distributions. *Sankhyā: The Indian Journal of Statistics (1933-1960)*, 19(1/2):23–26, 1958.
- [VDG96] Sara Van De Geer. Rates of convergence for the maximum likelihood estimator in mixture models. *Journal of Nonparametric Statistics*, 6(4):293–310, 1996.
- [vdVW96] Aad W. van der Vaart and Jon A. Wellner. *Weak Convergence and Empirical Processes: With Applications to Statistics*. Springer Series in Statistics. Springer, 1996.
- [Vil03] Cédric Villani. *Topics in Optimal Transportation*. American Mathematical Society, 2003.
- [Vil09] Cédric Villani. *Optimal transport : old and new*. Grundlehren der mathematischen Wissenschaften. Springer, Berlin, 2009.
- [VN98] D. L. Vandev and N. M. Neykov*. About regression estimators with high breakdown point. *Statistics*, 32(2):111–129, 1998.
- [YYC16] Chun Yu, Weixin Yao, and Kun Chen. A new method for robust mixture regression. *Canadian Journal of Statistics*, 45(1):77–94, 2016.

Index

The page number refers to the definition of the indexed topic.

- (a, b) -standard, 41
- $R(P_1, \dots, P_k)$, 144
- $R(\mathbf{t})$, 144
- $V(A)$, 126
- δ -covering number, 128
- δ -entropy number, 128
- δ -net, 128
- δ -packing number, 128
- δ -shattered, 130
- $\delta_{P,l}$, 30
- ϵ -covering number, 145
- ϵ -offset, 28
- \mathbf{c}^* , 82
- $\mathcal{P}_h(P)$, 141
- $\mathcal{P}_{\phi, \mathbf{c}, h}(P)$, 80
- $\mathcal{P}_{x,h}(P)$, 141
- $\mathcal{B}_\phi(\mathbf{c}, r)$, 80
- $\overline{\mathcal{M}}_h(P)$, 158
- $\partial_x f$, 160
- $c_{P,h}(v)$, 157
- h -trimmed k -optimal codebook, 82
- h -trimmed variation about \mathbf{c} , 78
- k -power distance, 142
- k -power distance-to-measure, 143
- k -variation given P , 78
- $m(P_{x,h})$, 141
- $r_{\phi,h}(\mathbf{c})$, 80
- $V_\phi^P(\mathbf{c})$, 80
- $V_{\phi,h}^P(\mathbf{c})$, 80
- $V_{\phi,k,h}(P)$, 81
- $V_{\phi,k}^P$, 81
- d_ϕ , 80
- $d_{P,h}$, 30
- $\text{dom}(f)$, 161
- bounded differences property, 124
- Bregman ball, 80
- Bregman divergence, 80
- Bregman divergence to the measure, 84
- closed, 161
- codebook, 77
- conjugate function, 161
- convex function of Legendre type, 161
- covering number, 145
- cumulant function, 97
- distance-to-measure, 30
- DTM-signature, 35
- effective domain, 161
- embedded, 36
- empirical k -PDTM, 147
- empirical distribution, 123
- extreme point, 180
- Gromov–Wasserstein distance, 35
- hard clustering, 77
- in general position, 159

- isomorphic, 34
- isomorphism, 34
- isotropic Gaussian model, 95
- label, 77
- likelihood, 93
- local mean, 141
- log-likelihood, 93
- loss, 78

- maximum likelihood estimator, 93
- metric entropy, 128
- metric-measure space, 34
- mixture model, 97
- mm-space, 34

- normalized mutual information, 88

- optimal codebook, 78
- optimal quantizer, 78

- parameter space, 97
- power distance, 142
- power function, 142
- proper, 161
- pseudo-distance, 30

- Rademacher average, 125
- reach, 32

- set of local distributions, 141
- set of local means, 141
- shattering dimension, 130
- signature, 35
- soft clustering, 77
- strongly convex, 80
- subgradient, 160
- subgradient set, 160
- sufficient statistic, 97
- supervised clustering, 77

- transport plans, 31
- trimmed k -means, 78
- trimmed k -variation, 78
- trimmed likelihood estimator, 93
- trimmed log-likelihood, 93
- trimming parameter, 103

- uniform measure, 32
- unsupervised clustering, 77

- variation about c given P , 78

- VC dimension, 126
- VC shatter coefficient, 126

- weighted Voronoi measures, 144

Titre : Vers une vision robuste de l'inférence géométrique.

Mots Clefs : Analyse géométrique des données, distance à la mesure, tests statistiques, partitionnement/quantification, inférence de support.

Résumé : Le volume de données disponibles est en perpétuelle expansion. Il est primordial de fournir des méthodes efficaces et robustes permettant d'en extraire des informations pertinentes. Nous nous focalisons sur des données pouvant être représentées sous la forme de nuages de points dans un certain espace muni d'une métrique, e.g. l'espace Euclidien \mathbb{R}^d , générées selon une certaine distribution. Parmi les questions naturelles que l'on peut se poser lorsque l'on a accès à des données, trois d'entre elles sont abordées dans cette thèse. La première concerne la *comparaison de deux ensembles de points*. Comment décider si deux nuages de points sont issus de formes ou de distributions similaires ? Nous construisons un test statistique permettant de décider si deux nuages de points sont issus de distributions égales (modulo un certain type de transformations e.g. symétries, translations, rotations...). La seconde question concerne la *décomposition d'un ensemble de points en plusieurs groupes*. Étant donné un nuage de points, comment faire des groupes pertinents ? Souvent, cela consiste à choisir un système de k représentants et à associer chaque point au représentant qui lui est le plus proche, en un sens à définir. Nous développons des méthodes adaptées à des données échantillonnées selon certains mélanges de k distributions, en présence de données aberrantes. Enfin, lorsque les données n'ont pas naturellement une structure en k groupes, par exemple, lorsqu'elles sont échantillonnées à proximité d'une sous-variété de \mathbb{R}^d , une question plus pertinente est de construire un système de k représentants, avec k grand, à partir duquel on puisse retrouver la sous-variété. Cette troisième question recouvre le problème de la *quantification* d'une part, et le problème de l'*approximation de la distance à un ensemble* d'autre part. Pour ce faire, nous introduisons et étudions une variante de la méthode des k -moyennes adaptée à la présence de données aberrantes dans le contexte de la quantification. Les réponses que nous apportons à ces trois questions dans cette thèse sont de deux types, théoriques et algorithmiques. Les méthodes proposées reposent sur des objets continus construits à partir de distributions et de sous-mesures. Des études statistiques permettent de mesurer la proximité entre les objets empiriques et les objets continus correspondants. Ces méthodes sont faciles à implémenter en pratique lorsque des nuages de points sont à disposition. L'outil principal utilisé dans cette thèse est la fonction distance à la mesure, introduite à l'origine pour adapter les méthodes d'analyse topologique des données à des nuages de points corrompus par des données aberrantes.

Title : Towards a Robust Vision of Geometric Inference.

Keys words : Geometric data analysis, distance-to-measure, statistical tests, clustering/quantization, support inference.

Abstract : It is primordial to establish effective and robust methods to extract pertinent information from datasets. We focus on datasets that can be represented as point clouds in some metric space, e.g. Euclidean space \mathbb{R}^d ; and that are generated according to some distribution. Of the natural questions that may arise when one has access to data, three are addressed in this thesis. The first question concerns the *comparison of two sets of points*. How to decide whether two datasets have been generated according to similar distributions ? We build a statistical test allowing to one to decide whether two point clouds have been generated from distributions that are equal (up to some rigid transformation e.g. symmetry, translation, rotation...). The second question is about the *decomposition of a set of points into clusters*. Given a point cloud, how does one make relevant clusters ? Often, it consists of selecting a set of k representatives, and associating every point to its closest representative (in some sense to be defined). We develop methods suited to data sampled according to some mixture of k distributions, possibly with outliers. Finally, when the data can not be grouped naturally into k clusters, e.g. when they are generated in a close neighborhood of some sub-manifold in \mathbb{R}^d , a more relevant question is the following. How to build a system of k representatives, with k large, from which it is possible to recover the sub-manifold? This last question is related to the problems of *quantization* and *compact set inference*. To address it, we introduce and study a modification of the k -means method adapted to the presence of outliers, in the context of quantization. The answers we bring in this thesis are of two types, theoretical and algorithmic. The methods we develop are based on continuous objects built from distributions and sub-measures. Statistical studies allow us to measure the proximity between the empirical objects and the continuous ones. These methods are easy to implement in practice, when samples of points are available. The main tool in this thesis is the function distance-to-measure, which was originally introduced to make topological data analysis work in the presence of outliers.

