



HAL
open science

Statistics for Topological Descriptors using optimal transport

Théo Lacombe

► **To cite this version:**

Théo Lacombe. Statistics for Topological Descriptors using optimal transport. Metric Geometry [math.MG]. Institut Polytechnique de Paris, 2020. English. NNT : 2020IPPAX036 . tel-02979251v2

HAL Id: tel-02979251

<https://hal.science/tel-02979251v2>

Submitted on 28 Oct 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



INSTITUT
POLYTECHNIQUE
DE PARIS

NNT : 2020IPPAX036

Thèse de doctorat



Statistiques sur les descripteurs topologiques à base de transport optimal

Thèse de doctorat de l'Institut Polytechnique de Paris
préparée à l'École polytechnique

École doctorale n°626 École Doctorale de l'Institut Polytechnique de Paris (ED IP
Paris)

Spécialité de doctorat : Mathématiques et Informatique

Thèse présentée et soutenue à Palaiseau, le 8 septembre 2020, par

THÉO LACOMBE

Composition du Jury :

Gabriel Peyré Directeur de Recherche CNRS, École Normale Supérieure (DMA)	Président
Peter Bubenik Professeur, Université de Floride (Département de mathématiques)	Rapporteur
François-Xavier Vialard Professeur, Université Gustave Eiffel (LIGM)	Rapporteur
Anthea Monod Professeur assistant, Imperial College London (Département de mathématiques)	Examineur
Sayan Mukherjee Professeur, Université Duke (Département de mathématiques)	Examineur
Steve Oudot Chargé de Recherche, Inria Saclay (Datashape)	Directeur de thèse
Marco Cuturi Professeur, Google Brain & ENSAE (CREST)	Co-directeur de thèse

REMERCIEMENTS

Je remercie tout d'abord les membres du jury qui ont chaleureusement accepté de participer à ma soutenance de thèse, Sayan Mukherjee, Anthea Monod, Gabriel Peyré, et particulièrement Peter Bubenik et François-Xavier Vialard qui ont rapporté ce manuscrit qui conclut trois années de recherche. J'en profite pour signaler que cette thèse a été rendue possible au moyen d'un financement AMX de l'École polytechnique.

Je remercie mes directeurs de thèse, Marco Cuturi et Steve Oudot, pour avoir guidé mes premiers pas de jeune chercheur, pour la confiance et la liberté qu'ils m'auront accordées en ayant toujours répondu présents dans les moments importants. Merci aux autres membres permanents de l'équipe DataShape, en particulier à Fred pour l'aide que tu m'as apportée, à Marc pour toutes les fois où tu as bien voulu écouter mes tergiversations scientifiques. Aux post-docs que j'ai vu passer, Miro, Hari, Jisu, et ce sacré Martin. Merci aux doctorants qui m'ont montré la voie, Claire, Jérémy, et le respecté docteur Mathieu qui m'a accueilli avec toute sa sympathie. À Nicolas pour m'avoir un jour parlé de la TDA. Aux joyeux compères qui sont arrivés avec moi, Raphaël et Vincent. Aux plus jeunes qui viennent mettre le chantier, Alex, Vadim, Étienne et Louis ; merci pour l'air frais que vous avez amené. Une pensée pour le bon Younes avec qui j'ai bien ri. Aux camarades de l'ENSAE, Boris et François-Pierre, votre compagnie quotidienne dans ce Slack désert était plus que bienvenue. À Christine, Stéphanie et Bahar pour leur disponibilité et leur sympathie. Toutes ces personnalités ont rythmé ma thèse au quotidien et leur présence se retrouve de près ou de loin dans mon travail ; j'en profite pour remercier une seconde fois ceux avec qui j'ai eu la chance de collaborer : Steve et Marco, Mathieu, Fred et Martin, et enfin Vincent - travailler avec toi fut un plaisir.

Un grand merci à tous ceux qui m'ont, à un moment où un autre, fait me sentir entouré. À mes amis de toujours, Jean-Baptiste, Arnaud et Vincent, pour leur présence inconditionnelle, même de l'autre côté du globe. Une pensée à mes amis de prépa, Léa, Rémy, Maxime, Paul, et bien d'autres,

avec qui je me suis formé aux mathématiques et à la khontrée. Aux amis de l'X, Benoît, Rémi, Victor, Pierre-Yves, Laure, le Bad', le Forum et tous les autres. La petite équipe TSN, Raymond, Ilyes et Julien, toujours là quand il fallait se changer les idées. À tous les membres de ma famille, à Audrey et Charline pour tous ces moments passés ensemble, à mes grands-parents Jeannine, Jean et Josette qui m'ont inspiré. À mes frères adorés Hugo et Dorian qui ont partagé mon enfance et qui tracent maintenant leurs chemins. À mes parents qui ont cru en moi depuis le début ; cette thèse, c'est aussi le fruit de votre soutien.

Enfin, merci à toi Lucie de continuer à illuminer mes journées à l'aube de cette treizième année qui, j'espère, se terminera au pays du soleil levant.

CONTENTS

1	Introduction	9
1.1	Introduction en français	9
1.2	Introduction in English	20
2	Background	31
2.1	Topological Data Analysis	31
2.2	Optimal Transport	45
2.3	Notations	59
I	Theory	61
3	Persistence diagrams and measures, an Optimal Transport viewpoint	63
3.1	General properties	63
3.2	Persistence measures in the finite setting	69
3.3	The bottleneck distance	72
3.4	Duality	76
3.5	Proofs	79
4	Fréchet means in the space of persistence measures	91
4.1	Fréchet means in the finite case	92
4.2	Existence and consistency	93
4.3	Fréchet means of persistence diagrams	96
4.4	Proofs	97
II	Applications in statistics and learning	107
5	Fast estimation of Fréchet means of persistence diagrams	109

5.1	Preliminary remarks and problem formulation	110
5.2	A Lagrangian approach	112
5.3	Entropic Regularization	118
6	Linear representations	133
6.1	Continuity of linear representations	133
6.2	Learning representations using PersLay	137
7	Complementary examples	147
7.1	Expected persistence diagrams	148
7.2	Quantization of persistence diagrams	157
7.3	Shift-invariant distance	167
	Conclusion	178
	Bibliography	183
	Index	199
A	Homology theory	203
A.1	Homology theory	203
A.2	Filtrations and persistence modules	207
A.3	Persistence diagrams	212
B	Elements of measure theory	221
C	Elements of metric geometry	223

LIST OF FIGURES

1.1	Mélange de Gaussiennes	10
1.2	Préservation de la topologie par certaines transformations . . .	11
1.3	Topologie a différentes échelles	12
1.4	Diagrammes et distance de bottleneck	14
1.5	Stabilité de la bottleneck	15
1.6	Gaussian mixture	21
1.7	Homotopy transformation	22
1.8	Scale variant topology	23
1.9	Diagrams and bottleneck distance	25
1.10	Stability of the bottleneck distance	26
2.1	Partial matching between two diagrams	33
2.2	Stability of the bottleneck distance (sketch)	35
2.3	Instability of cardinality of diagrams in sampling process	40
2.4	Cauchy sequence in the space of persistence diagrams	41
2.5	Fréchet mean of three persistence diagrams	43
2.6	Optimal transport problem in 1D	46
2.7	Transport map vs transport plan	48
2.8	Wasserstein barycenter (example)	51
2.9	Induced transport map in partial optimal transport	54
3.1	Illustration of differences between OT_p , OT_∞ , and vague con- vergences.	75
4.1	Global picture of the proof of Proposition 4.4	99
4.2	Partition used in the proof of Lemma 4.7.	100
5.1	Example of output of Algorithm 1	114
5.2	Local minimizers of the B-Munkres algorithm	116
5.3	Error control in regularized OT	121
5.4	Convolution in Sinkhorn algorithm (sketch)	124

5.5	Parallelism in Sinkhorn algorithm (sketch)	125
5.6	Regularized Barycenter estimation for persistence diagrams	126
5.7	Fréchet mean vs arithmetic mean	128
5.8	Average running times for B-Munkres and Sinkhorn barycenters for persistence diagrams	129
5.9	Qualitative comparison of B-Munkres and our Algorithm.	129
5.10	k -means in the space of persistence diagrams	130
6.1	Some common linear representations of persistence diagrams	134
6.2	Orbits of dynamical systems	141
6.3	PERSLAY (sketch)	144
6.4	Learning the weight function	145
7.1	Expected persistence diagram (example)	149
7.2	Stability of expected persistence diagrams	150
7.3	Quantization of persistence diagrams	167
7.4	Čech diagram in log scale	169
7.5	Persistence transform	170
7.6	Non convexity of diagram distance up to translation	171
7.7	Shift with a single point	173
7.8	Shift-invariant distance	179
A.1	An example of simplicial complex	204
A.2	Singular simplex (sketch)	208
A.3	Čech and Rips complexes	209
A.4	Algebraic pipeline to build persistence modules (sketch)	211
A.5	Filtration and persistence module (example)	212
A.6	Persistence homology pipeline (example)	217
A.7	Interleaving between persistence module (sketch)	218

CHAPTER 1

INTRODUCTION

1.1 Introduction en français

1.1.1 Analyse des données, statistiques et géométrie

L'analyse des données. L'analyse des données est devenue, au cours de la dernière décennie, un des domaines de recherche en mathématiques appliquées et en informatique les plus actifs et prolifiques. L'évolution des moyens d'acquisition et des capacités de stockage en tout genre a permis de constituer des collections de données colossales en tout genre : images [KH⁺09], formes 3D [CFG⁺15], sons [BMEWL11], données médicales et biologiques [CHF12], réseaux sociaux [YV15, KKM⁺16, PWZ⁺17]... Le but de l'analyse des données est de comprendre et de valoriser ces nouvelles informations disponibles.

Pour cela, le paradigme le plus courant de nos jours est de recourir à "l'apprentissage automatique", dont l'idée générale est la suivante : proposer un algorithme capable d'extraire une information utile d'un jeu de données, permettant par exemple de réaliser de regrouper les données par similarité (*clustering*), d'attribuer des labels à de nouvelles observations (*classification*), etc.

Statistiques et géométrie. Évidemment, il n'existe pas d'algorithme tout-puissant, qui serait capable de prédire sans erreur dans n'importe quel contexte. Si les observations et les labels ne sont pas ou peu liés, ou si le nombre d'observations est trop faible, même le meilleur algorithme possible ne saura pas se montrer utile lorsqu'il s'agira d'extrapoler son apprentissage à de nouvelles observations.




Figure 1.1: Nuages de 400 points tirés selon un mélange de quatre gaussiennes, représentant quatre classes. Sur la figure de gauche, les gaussiennes ont une faible variance, et il est facile de séparer les différentes classes. Sur la figure de droite, la séparation des classes est plus difficile.

La difficulté à résoudre un problème d'apprentissage, par exemple sa dépendance au nombre d'observations, peut se mesurer au moyen d'outils statistiques. Dans un problème de classification ou de clustering par exemple, plus les observations d'une même classe (c-à-d. partageant un même label) sont concentrées autour de leur moyenne (faible *variance*) et plus les classes sont distinctes, plus le problème sera résolu efficacement (voir Figure 1.1) De manière générale, il est important de comprendre comment nos observations sont *distribuées*. L'approche généralement adoptée en statistique est la suivante : nos observations sont *indépendantes* et proviennent d'une distribution de probabilité sous-jacente (inconnue). Connaître parfaitement cette distribution permettrait de comprendre le comportement des algorithmes d'apprentissage et d'optimiser leurs performances sur ce jeu de données (sans pour autant arriver nécessairement à un score parfait). En pratique, l'idée est donc d'estimer certaines propriétés de cette distribution sous-jacente à partir de notre *échantillon* d'observations.

L'exemple le plus élémentaire, et qui va nous amener à des considérations géométriques, est l'estimation de l'*espérance* de la loi sous-jacente, c'est-à-dire la valeur moyenne d'une observation tirée aléatoirement. Ce problème est en général très simple quand on observe des nombres réels : si on observe $(x_1 \dots x_n)$ n nombres tirés indépendamment selon une loi μ , un estimateur de l'espérance de μ est simplement la moyenne arithmétique des $(x_i)_i$, c'est-à-dire $\frac{1}{n} \sum_{i=1}^n x_i$. Cependant, dans les problèmes d'apprentissage modernes,




Figure 1.2: À gauche, un tore dont on identifie les propriétés topologiques. Ensuite, un cercle, deux transformations qui en préservent la topologie, et deux transformations qui ne les préservent pas (déchirement et recollement respectivement).

la géométrie sous-jacente à nos observations est généralement complexe et inconnue : graphes, images, molécules... autant de cas où la notion de moyenne, et *a fortiori* des descripteurs statistiques plus sophistiqués, est difficile à proprement définir et estimer.

1.1.2 Descripteurs topologiques

Dans cette thèse, nous nous intéresserons à un type d'observations en particulier, les *diagrammes de persistance*, qui proviennent de *l'analyse topologique des données*. L'étude de la géométrie sous-jacente à ces objets, tout comme la conception et le calcul de descripteurs statistiques, sont des sujets de recherche actifs.

Topologie. Brièvement, la notion de topologie renvoie à celle de “forme”. Étant donné un objet, par exemple une forme 3D, combien celui-ci possède-t-il de composantes connexes ? Peut-on identifier des boucles caractéristiques à sa surface ? Possède-t-il des cavités ? Le tore représenté en Figure 1.2 possède par exemple une composante connexe, deux boucles (en rouge), et une cavité (en bleu). Le cercle, lui, possède une composante connexe et une boucle.

Dans un cadre plus général, la topologie peut se comprendre comme les propriétés d'un objet qui sont préservées par des transformations continues, sans déchirement ni recollement, voir Figure 1.2. Par exemple, lorsqu'on applique une telle transformation à un cercle, si la géométrie de l'objet peut varier, sa topologie reste inchangée. D'un point de vue formel, les propriétés topologiques d'un objet sont décrites par ses *groupes d'homologie*. Les groupes d'homologie d'un espace topologique (objet) sont décrit par une suite de groupe abélien $(H_k)_k$ ($k = 0$ correspond aux propriétés topologiques de dimension 0, c'est-à-dire les composantes connexes; $k = 1$ correspond aux boucles, $k = 2$ aux cavités, etc.) dont les générateurs indentifient les




Figure 1.3: Illustration du phénomène d'évolution de la topologie à différentes échelles.

propriétés topologiques “indépendantes” de l'espace. Par exemple, pour un tore, H_1 a deux générateurs. Ceux-ci correspondent intuitivement aux deux boucles rouges sur la Figure 1.2: toutes les autres boucles que l'on peut dessiner sur le tore s'écrivent - en un sens - comme une combinaison linéaire de ces deux boucles. Le lecteur intéressé peut consulter l'Annexe A ou [Mun84] pour une présentation plus détaillée de la théorie de l'homologie.

Descripteurs multi-échelles de la topologie. En pratique néanmoins, l'usage de la topologie en apprentissage automatique n'est pas immédiat. La plupart des méthodes d'acquisition des données vont produire des objets sous forme de *nuage de points*, dont la topologie intrinsèque est excessivement simple. Les groupes d'homologie d'un nuage de points $(H_k)_k$ sont triviaux dès que $k \geq 1$, tandis que H_0 est le groupe abélien libre engendré par N éléments, où N désigne la cardinalité du nuage de points. L'homologie seule n'est pas capable de refléter la structure sous-jacente de tels objets.

L'idée fondamentale développée dans les années 2000 (quoique les fondements peuvent être tracés tout au long du XX^e siècle) est de regarder la topologie d'un espace topologique à *différentes échelles*, et de regarder quelles sont les propriétés topologiques qui *persistent* à travers celles-ci. Considérons par exemple le nuage de points représenté à gauche sur la Figure 1.3. Initialement, il ne s'agit (topologiquement) que d'une collection de composantes connexes indépendantes. Il apparaît cependant naturel d'identifier trois boucles, de tailles différentes. Regarder la topologie de cet objet à différentes échelles va nous permettre de détecter lesdites boucles. Pour introduire cette notion d'échelle, une idée est de faire grossir des boules centrées en chacun des points du nuage. Le paramètre d'échelle $t \geq 0$ correspond ici au rayon des boules. Ainsi, à partir du nuage de points X au paramètre $t = 0$, on construit une famille d'objets $(X_t)_{t \geq 0}$. On parle alors de *filtration*. Pour certaines *valeurs critiques* du paramètre t , voir Figure 1.3, la topologie¹ change : des boucles apparaissent (par exemple à

¹Ici, par souci de simplicité, on s'intéresse uniquement aux boucles.

$t = 1$ ou $t = 2$ sur la figure) ou disparaissent (lorsqu’elles sont complètement “remplies”, $t = 2$, $t = 3$, $t = 5$ sur la figure).

Ce sont ces valeurs critiques, appelées *temps de naissance* et de *mort* des propriétés topologiques, qui seront enregistrées dans un descripteur topologique : le *diagramme de persistance*. Un diagramme de persistance se présente comme une collection de points dans le plan. La présence d’un point de coordonnées (b, d) dans le diagramme va indiquer qu’une propriété topologique (composante connexe, boucle, cavité...) est apparue à l’échelle $t = b$ et a disparu à l’échelle $t = d$. Ainsi, un point proche de la diagonale, c’est-à-dire tel que $d \simeq b$, représente une propriété topologique qui est apparue et a presque immédiatement disparu : cette propriété a peu *persisté* à travers les échelles. À l’inverse, un point loin de la diagonale va représenter une composante topologique présente durant un large intervalle d’échelles, généralement considérée comme plus significative. Formellement, un diagramme de persistance est décrit comme un multi-ensemble de points² supporté sur le demi-plan

$$\Omega := \{(b, d) \in \mathbb{R}^2, d > b\}$$

ou, de façon équivalente, comme une mesure de Radon μ s’écrivant

$$\mu := \sum_{x \in X} n_x \delta_x,$$

où $X \subset \Omega$ est localement fini, n_x un entier et δ_x désigne la masse de Dirac en $x \in X$.³

Bien entendu, le cadre d’application de l’analyse topologique des données ne se résume pas aux nuages de points et à des boules qui grossissent ; la théorie générale est présentée dans la Section 2.1. À ce stade, l’essentiel est de retenir qu’il est donc possible de transformer une collection d’observations complexes en une collection de diagrammes de persistance, à partir de laquelle on peut envisager de produire une analyse statistique ou de réaliser une tâche d’apprentissage automatique.

²Un ensemble dans lequel les points peuvent être répétés

³Dans leur définition la plus générale, les diagrammes de persistance peuvent avoir des points avec des coordonnées infinies ($b = -\infty$ or $d = +\infty$) qui appartiennent donc au demi-plan étendu; ou des points sur la diagonale ($b = d$). Nous ignorons ces points dans cette thèse, voir les Remarques 2.2 et 2.3.




Figure 1.4: Deux diagrammes de persistance. La distance de bottleneck entre ces deux diagrammes est, par définition, la longueur de la plus longue arête dessinée.

1.1.3 Diagramme de persistance et apprentissage: limitations et enjeux

Les diagrammes de persistance portent une information riche et en général complémentaire aux autres méthodes classiques d'analyse de données. Malheureusement, leur incorporation dans les outils modernes d'apprentissage, tels que les réseaux de neurones, n'est absolument pas évidente.

Distance entre diagrammes de persistance. Pour comprendre cela, il faut dans un premier temps souligner qu'il est possible de mesurer une distance entre deux diagrammes de persistance, ce qui permet donc de comparer les deux objets initiaux d'un point de vue topologique. La distance de référence entre diagrammes est appelée *distance de bottleneck*, et se calcule de la façon suivante. Pour deux diagrammes X et Y , notons $(x_1 \dots x_n)$ et $(y_1 \dots y_m)$ leurs points respectifs (notons qu'on n'a pas nécessairement $n = m$). La première étape est de chercher à *transporter*⁴ chaque point x_i du diagramme X vers un point y_j de Y ou, éventuellement, vers sa projection orthogonale sur la diagonale (voir Figure 1.4). Les points y_j de Y qui ne sont pas atteints par un point x_i de X sont eux aussi transportés sur leur projection sur la diagonale. Nous imposons de plus que ce transport soit *bijectif* : chaque point x_i est envoyé - au plus - sur un point y_j , et chaque y_j doit être atteint par - au plus - un seul x_i . En notant (x_i, y_j) quand x_i est transporté sur y_j , et $(x_i, \partial\Omega)$ (resp. $(\partial\Omega, y_j)$) lorsque x_i (resp. y_j) est transporté sur la diagonale, un *transport partiel* entre X et Y est décrit par une liste $\dots (x_i, y_j)_{ij} \dots (x_i, \partial\Omega) \dots (\partial\Omega, y_j) \dots$, où chaque point x_i de X et chaque point y_j de Y apparaît exactement une fois.

⁴On parle plus fréquemment "d'appariement", mais nous préférons la terminologie de "transport" pour des raisons qui deviendront évidentes par la suite.




Figure 1.5: Illustration de la stabilité des diagrammes de persistance en bottleneck, et instabilité du nombre de points.

Bien entendu, il existe une multitude de façons de transporter un diagramme X sur un autre diagramme Y . Pour définir la distance entre deux diagrammes, nous allons chercher un transport partiel *optimal* au sens suivant : le *coût* d'un transport partiel est la plus grande distance parcourue entre deux points du transport. Un transport partiel optimal est un transport partiel dont le coût est minimal parmi tous les transports partiels entre X et Y possibles. Le coût minimal ainsi réalisé est, par définition, la distance de bottleneck. Formellement, si on note $\Gamma(X, Y)$ l'ensemble des transports entre deux diagrammes X et Y , on pose

$$\mathbf{d}_\infty(X, Y) := \min_{\gamma \in \Gamma(X, Y)} \max_{(x, y) \in \gamma} \|x - y\|.$$

Le choix de cette distance n'est pas arbitraire et est motivé par des considérations algébriques profondes qui sont détaillées dans la Section 2.1. Cette distance jouit de la propriété d'être *stable*, au sens où des objets initialement proches auront systématiquement des diagrammes proches en distance de bottleneck (Figure 1.5).

Limitations. Étant donnée une collection de diagrammes, grâce à la distance de bottleneck, il est possible de calculer des distances entre chaque paire de diagrammes. L'utilisation de certains algorithmes d'apprentissage basiques, comme celui du plus proche voisin, ou par exemple le *multi-dimensional scaling* (MDS), est alors possible, car ceux-ci requièrent seulement de savoir calculer des distances entre nos observations (ici les diagrammes).

Néanmoins, la plupart des algorithmes modernes d'apprentissage demandent beaucoup plus de structure, qui va faire défaut dans le cadre des diagrammes de persistance. En effet, *l'espace* des diagrammes de persistance n'est pas *linéaire*, ce qui veut essentiellement dire qu'il n'est pas possible de donner un sens⁵ à la somme de deux diagrammes de persistance, ou à la multiplication d'un diagramme par un nombre réel. Il n'est donc pas possible d'incorporer naïvement les diagrammes de persistance dans les machineries d'apprentissage modernes, et leur utilisation dans la pratique en est donc compromise.

Une solution de remplacement va consister à *plonger* l'espace des diagrammes dans un espace linéaire, sur lequel il sera donc possible d'effectuer des tâches d'apprentissage aisément. Cela signifie qu'on cherche une application Φ qui transforme les diagrammes en *vecteurs*, c'est-à-dire un élément dans un espace linéaire. Une telle application est appelée une *vectorisation*. Par exemple, une vectorisation naïve pourrait consister à associer à chaque diagramme X son nombre de points $\Phi(X)$, plongeant l'espace des diagrammes dans l'espace des nombres réels, bien entendu linéaire. Cette vectorisation est très décevante car deux diagrammes très différents (au sens de la distance de bottleneck) peuvent avoir le même nombre de points, mais aussi deux diagrammes très proches peuvent avoir un nombre de points très différent (comme en Figure 1.5). Autrement dit, $\Phi(X)$ n'a pas grand chose à voir avec X , ce qui est profondément regrettable.

Bien entendu, de nombreuses vectorisations bien plus satisfaisantes ont été proposées, avec d'importants succès dans les applications [Bub15, AEK⁺17, CCO17]. Néanmoins, l'usage d'une vectorisation soulève quelques interrogations importantes :

- Le choix de la vectorisation est arbitraire. Parmi le large catalogue de vectorisations disponibles, il n'existe pas à ce jour d'heuristique pour choisir laquelle sera adaptée à une tâche d'apprentissage donnée.
- Il semble impossible⁶ d'obtenir, en toute généralité, un plongement qui soit *bi-stable*, c'est à dire pour lequel on aurait, pour toute paire de diagramme X, Y , et deux constantes $0 < A \leq B < \infty$,

$$A\|\Phi(X) - \Phi(Y)\| \leq \mathbf{d}_\infty(X, Y) \leq B\|\Phi(X) - \Phi(Y)\|,$$

ce qui aurait le mérite d'assurer que des diagrammes proches ont des représentations proches, et inversement.

⁵Qui serait compatible avec la distance de bottleneck

⁶Ce problème est encore ouvert, mais tous les résultats actuels semblent aller dans ce sens [CB19, BW20, Wag19]

- Enfin, il faut rappeler que les diagrammes de persistance ont été construits en partie pour leur interprétabilité. Recourir à une vectorisation, et réaliser des calculs dans l'espace linéaire associé, retire cette interprétabilité. S'il est possible pour une collection de diagrammes $X_1 \dots X_n$ de calculer facilement une moyenne dans l'espace linéaire en posant $M = \frac{1}{n} \sum_{i=1}^n \Phi(X_i)$, il est important de comprendre que cette quantité n'a, *a priori*, rien à voir avec une notion de moyenne dans l'espace des diagrammes de persistance. Cela vaut plus généralement pour toutes les opérations réalisées dans l'espace linéaire (qui conduiraient par exemple à du clustering, etc.).

Enjeux. Si l'étude des vectorisations de l'espace des diagrammes de persistance reste un domaine de recherche pertinent, nous proposons ici de travailler *directement dans l'espace des diagrammes de persistance*. L'objectif est donc de répondre à la question

“Comment peut-on faire des statistiques dans l'espace des diagrammes de persistance ?”

Pour cela, il est important de comprendre d'abord quelles sont les propriétés géométriques de cet espace, pour lequel nous ne disposons *a priori* que d'une distance pour en étudier la structure. Cela doit permettre ensuite de proposer un cadre théorique solide dans lequel des outils statistiques sont proprement définis. Enfin, il est important que ces outils statistiques puissent être calculés et utilisés, au moins de façon approchée, en pratique.

1.1.4 Contributions et organisation du manuscrit

Afin de proposer des éléments de réponse à cette problématique, nous nous appuyons sur un autre domaine des mathématiques appliquées : la théorie du *transport optimal*, dont la présentation est faite dans la Section 2.2. Il s'agit peut-être du point essentiel de ce manuscrit : l'établissement d'un lien formel entre l'espace des diagrammes de persistance et les modèles utilisés en transport optimal. Cette connexion entre ces deux champs mathématiques est extrêmement prolifique : le transport optimal est un domaine très développé, tant dans sa théorie que dans les applications en statistiques et en apprentissage ; et nombre de ses outils s'adaptent à l'étude des diagrammes de persistance. Cela nous permettra d'obtenir de nouveaux résultats théoriques relatifs à l'utilisation des diagrammes de persistance en statistiques et en apprentissage, mais aussi de proposer divers algorithmes et outils numériques qui permettent de résoudre des problèmes difficiles

en persistance, comme l’estimation de moyennes (dites “de Fréchet”), la quantification, l’apprentissage de vectorisation, entre autres.

De façon générale, et sauf mention explicite du contraire, nous adoptons la convention suivante : les résultats (théorèmes, propositions, etc.) pour lesquels nous proposons une preuve sont le fruit de mon travail ou d’un travail joint avec mes collaborateurs : Mathieu Carrière, Frédéric Chazal, Marco Cuturi, Vincent Divol, Yuichi Ike, Steve Oudot, Martin Royer, et Yuhei Umeda.

Plan et lien avec les travaux réalisés durant la thèse.

Chapitre 2: Préliminaires. Ce chapitre présente les deux domaines impliqués dans ce travail: l’analyse topologique des données (Section 2.1) et le transport optimal (Section 2.2). L’objectif est de présenter rapidement l’état de l’art relatif aux différents problèmes abordés dans ce manuscrit. Notons néanmoins que la construction algébrique des diagrammes de persistance, développée dans les sous-sections A.1 et A.2, est indépendante du reste du manuscrit.

Les contributions sont présentées en deux parties, en fonction de leur nature théorique (propriétés géométriques générales, etc.) ou appliquée (pour faire simple, conduisant à une implémentation).

Partie I — Théorie.

Chapitre 3: Un formalisme issu du transport optimal pour les diagrammes de persistance. Ce chapitre présente et développe les fondements théoriques de ce manuscrit. Son contenu repose essentiellement sur la section 3 de l’article [DL19], en révisions mineures au *Journal of Applied and Computational Topology*. Nous y présentons comment les métriques utilisées habituellement pour comparer les diagrammes de persistance peuvent se reformuler comme des problèmes de transport optimal *partiels*, et les nombreux résultats théoriques qui en découlent.

Chapitre 4: Moyenne de Fréchet pour les diagrammes de persistance: aspects théoriques. Ce chapitre est consacré à l’étude des moyennes de Fréchet (ou barycentres) pour les diagrammes de persistance. Il repose sur la section 4 de [DL19]. On y prouve notamment un résultat d’existence très général, et on établit un lien fort entre ces moyennes et les fameux “barycentres de Wasserstein”, largement étudiés dans la littérature du transport optimal traditionnel.

Partie II — Applications.

Chapitre 5: Algorithmes efficaces pour l’estimation des moyennes de Fréchet des diagrammes de persistance. On propose ici un algorithme pour approcher les moyennes entre diagrammes; particulièrement efficace pour traiter les problèmes à grande échelle. Cette approche a été publiée dans les annales de la conférence internationale *Neural Information Processing Systems, 2018*, voir [LCO18].

Chapitre 6: Représentations linéaires des diagrammes de persistance. Ce chapitre se consacre à l’étude des représentations linéaires de diagramme de persistance, pour lesquelles on propose une caractérisation exhaustive. Une fois encore, les résultats théoriques (Section 6.1) sont issus de [DL19, §5.1]. On propose ensuite une application de ce résultat en apprentissage automatique en introduisant PERSLAY, une couche pour les réseaux de neurones⁷ spécifiquement élaborée pour apprendre des représentations adaptées à une tâche donnée. Ce travail a été publié dans les annales de la conférence *Artificial Intelligence and Statistics, 2020*, voir [CCI+20].

Chapitre 7: Exemples complémentaires. Ce dernier chapitre regroupe diverses applications et algorithmes qui ont été développés afin de résoudre divers problèmes relatifs aux diagrammes de persistance. La Section 7.1 démontre l’intérêt du formalisme théorique développé dans le Chapitre 3 pour étudier les diagrammes de persistance dans un contexte aléatoire. On y trouve des résultats de convergence et de stabilité pour des analogues aléatoires des diagrammes de persistance. La Section 7.2 étudie la *quantisation* des diagrammes de persistance. Enfin, la Section 7.3 propose une façon simple et efficace pour estimer des distances entre diagrammes *a transformation près*. Ces résultats n’ont pas encore été publiés, mais les algorithmes correspondant ont, ou vont être, intégrés à la librairie d’analyse topologique des données **Gudhi** [GUD15]; ce manuscrit de thèse est une bonne occasion d’en présenter les rouages.

Code et contributions à la librairie **gudhi.** La plupart des méthodes présentées en Partie II ont été, ou vont être, incorporées à la librairie **Gudhi**.

⁷disponible publiquement et incorporée à la librairie **Gudhi**.

1.2 Introduction in English

1.2.1 Data Analysis, Statistics, and Geometry

Data Analysis. Data analysis has become, during the last decade, one of the most active research areas in applied mathematics and computer science. Progress in data gathering and storage has led to very large datasets of various types: images [KH⁺09], 3D shapes [CFG⁺15], music [BMEWL11], medical and biological data [CHF12], social networks [YV15, KKM⁺16, PWZ⁺17]... Data analysis aims to understand and add value to such newly available information.

To do so, the most standard paradigm nowadays is to make use of “machine learning”. It aims at designing algorithms that will be able to extract useful information from a given set of observations; allowing for instance to regroup data by similarity (*clustering*), to assign labels to new observations based on the labels of the *training* data (*classification*), etc.

Statistics and geometry. Of course, there is no omnipotent algorithm that would be able to predict with no error in any context. If observations and labels are not correlated or if the number of observations is too low, even the best possible algorithm will not be of any use when it comes to generalizing its learning to new sets of observations.

The difficulty to solve a given learning problem, for instance its dependence on the number of observations, can be measured with statistical tools. In a classification or clustering problem for instance, the more the observations sharing the same label are concentrated around their mean value (low *variance*) and the more the different classes are separated from each other, the easier the problem will be to solve (see Figure 1.6). In general, it is important to understand how our observations are *distributed*. The general approach taken in statistical learning is to assume that our observations are *independent* and come from an (unknown) underlying probability distribution. Perfect knowledge of this distribution would allow us to understand the behavior of learning algorithms and to optimize their performances on a given dataset (without necessarily reaching a perfect score). In practice, it is thus useful to infer some properties of the underlying probability distribution from our *sample* of observations.

The most basic example, which will lead to geometric considerations, is the estimation of the *expectation* of the underlying law, that is the average value of a randomly sampled observation. When the observations are real numbers, this problem is pretty simple: if we observe $(x_1 \dots x_n)$ n numbers independently sampled from a law μ , an estimation of the expectation of




Figure 1.6: Point clouds (400 points) sampled from a mixture of four Gaussian distributions, representing four classes. On the left subfigure, the Gaussians distributions have a low variance, it is thus easy to separate the different classes. On the right subfigure, class separation is harder to perform.

μ is simply given by the arithmetic mean of the $(x_i)_i$ s, that is $\frac{1}{n} \sum_{i=1}^n x_i$. However, in modern machine learning problems, the underlying geometry of the observations is often complex and unknown: think of graphs, images, molecular structures... These are cases where the notion of mean, and *a fortiori* more sophisticated descriptors, is hard to clearly define and to estimate.

1.2.2 Topological Descriptors

In this thesis, we will focus on a specific type of observations: the *persistence diagrams*, which come from *topological data analysis*. Understanding the underlying geometry of these objects is an active research area, as is the design and the computation of statistical descriptors built on top of persistence diagrams.

Topology. Roughly speaking, the term “topology” is related to the notion of “shape”. Considering a given object, such as a 3D shape, how many connected components does it have? Can we identify characteristic loops? Does it have cavities? For instance, the torus represented in Figure 1.7 has one connected component, two loops (red), and a cavity (blue). In contrast, a circle has a single connected component and a single loop.




Figure 1.7: On the left, a torus and its topological properties. Then, a circle, two transformations that preserve its topology, and two transformations that do not (ripping and gluing respectively).

More generally, topology can be described as the set of properties of an object that do not change when we apply transformations without ripping nor gluing (see Figure 1.7). For instance, when such a transformation is applied to a circle, the topology (presence of one loop and one connected component) is unchanged, although the geometry is. Formally, the topological properties of an object are described by *homology groups*. Homology groups of a topological space (object) consist of a sequence of abelian groups $(H_k)_{k \geq 0}$ ($k = 0$ accounts for 0-dimensional topological properties of the space, that is connected components; $k = 1$ accounts for loops, $k = 2$ for cavities, and so on) whose generators identify “independent” topological features of the space. For instance, in the context of a torus (Figure 1.7), H_1 has two generators. These would correspond intuitively to the two red loops: any other loop on the torus is—in some sense—a linear combination of these two loops. We refer the interested reader to Appendix A or [Mun84] for a detailed presentation of homology theory.

Multi-scale topological descriptors. In practice however, the use of topology in machine learning is not straightforward. Most data acquisition methods are likely to produce objects with an extremely simple topology, such as *point clouds*. The homology groups of a point cloud $(H_k)_k$ are trivial for $k \geq 1$, while H_0 is the free abelian group with N generators, N being the cardinality of the point cloud. Homology alone is not able to reflect the underlying structure of such objects.

The core idea, whose foundations can be traced back to Morse’s work [Mor40] and have been developed in [Fro92, Rob99, CFP01, ZC05], is to consider the topology of a space *at different scales*, and to look for topological properties that *persist* through scales. Let us consider as an example the point cloud represented on the left of Figure 1.8. Initially, it is (from a topological perspective) a collection of independent connected components. It is however intuitive to identify three loops of different sizes. Recording the topology of this object at different scales will allow us to detect these




Figure 1.8: (Top) Topology of a point cloud at different scales for the Čech filtration. (Bot) The persistence diagram of the point cloud (restricted to loops).

loops. To introduce this notion of scale, an idea is to put balls centered on each point of the point cloud and to let them grow. The scale parameter $t \geq 0$ corresponds here to the radius of the balls. Therefore, starting with a point cloud X at $t = 0$, we build an increasing sequence of topological spaces $(X_t)_{t \geq 0}$, called a (Čech) *filtration*. For some *critical values* of t (see Figure 1.8) the topology⁸ changes: loops appear (at $t = 1, t = 1.5, t = 2$ on the figure) or disappear (when they get “filled in”, at $t = 2, t = 3, t = 5$ on the figure).

These critical values, called *birth times* and *death times*, are recorded in a topological descriptor: the *persistence diagram*, which is a collection of points in the plane. Each point of coordinate (b, d) in the diagram will account for a topological property (connected component, loop, caveat...) that appeared at scale $t = b$ and disappeared at scale $t = d$. Therefore, a point close to the diagonal (that is, such that $b \simeq d$) represents a topological property that appeared then almost instantly disappeared: this feature did not *persist* much. In contrast, a point away from the diagonal represents a topological feature that was recorded through a large interval of scales, generally considered as being more significant. Formally, a persistence diagram

⁸Here, for the sake of simplicity, we only consider loops.

can be described as a multi-set of points⁹ supported on a half-plane

$$\Omega := \{(b, d) \in \mathbb{R}^2, d > b\}$$

or, equivalently, as a *point measure*, that is a Radon measure μ of the form

$$\mu = \sum_{x \in X} n_x \delta_x,$$

where X is a locally finite subset of Ω , n_x is an integer and δ_x denotes the Dirac mass located at $x \in X$.¹⁰

Of course, the situations in which topological data analysis can be applied are not reduced to point clouds and growing balls: the general theory is detailed in Appendix A. The important point is that one can transform a set of complex observations into a collection of persistence diagrams, from which we might consider producing a statistical analysis or performing a machine learning task.

1.2.3 Persistence diagrams and learning: challenges and limitations

In practical applications, incorporating persistence diagrams in modern machine learning toolboxes, such as neural networks, is unfortunately not straightforward.

Distance between persistence diagrams. In order to understand this, we must first mention that one can compute a distance between two persistence diagrams, which thus allows us to compare two objects adopting a topological viewpoint. The standard distance used to compare diagrams is called the *bottleneck distance*, and can be computed in the following way. For two diagrams X and Y , let $(x_1 \dots x_n)$ and $(y_1 \dots y_m)$ denote their respective points (note that we do not necessarily have $n = m$). The first step is to *transport*¹¹ each point x_i in the diagram X to a point y_j in Y or, possibly, to its orthogonal projection onto the diagonal (see Figure 1.9). The points y_j in Y which are not reached by a point x_i in X are also transported

⁹A set where points can be repeated.

¹⁰In the greatest generality, persistence diagrams can have points with infinite coordinates ($b = -\infty$ or $d = +\infty$) thus belonging to the extended half-plane; or points supported exactly on the diagonal ($b = d$). We will not consider such points in this thesis; this is discussed in Remark 2.2 and Remark 2.3.

¹¹It is more standard to use the wording “matching”, but we will prefer “transport” for reasons that will become obvious later on.




Figure 1.9: Two persistence diagrams. The bottleneck distance between these two diagrams is given, by definition, by the length of the longest edge drawn here.

on their respective projections onto the diagonal. We also enforce that each point x_i is transported to—at most—a single point y_j , and each y_j must be reached by—at most—a single x_i . Denote by (x_i, y_j) the fact that x_i is transported on y_j , and $(x_i, \partial\Omega)$ (resp. $(\partial\Omega, y_j)$) when x_i (resp. y_j) is transported to the diagonal. A *partial transport* between X and Y is given by a list $\dots(x_i, y_j)_{ij} \dots (x_i, \partial\Omega) \dots (\partial\Omega, y_j) \dots$, where each x_i, y_j must appear exactly once.

Of course, there exist many ways to transport a diagram X onto another diagram Y . In order to define a distance, one seeks for an *optimal* partial transport in the following sense: the *cost* of a partial transport is given by the longest distance traveled when performing the transport. A partial transport is optimal if it has minimum cost among all the possible partial transport between x and Y . The bottleneck distance is then defined as the minimal cost achieved by this optimal transport. Formally, let $\Gamma(X, Y)$ denote the set of all possible partial transports between X and Y , and define the bottleneck distance as

$$\mathbf{d}_\infty(X, Y) := \min_{\gamma \in \Gamma(X, Y)} \max_{(x, y) \in \gamma} \|x - y\|.$$

The choice of such a distance is motivated by algebraic considerations detailed in Appendix A and Section 2.1. This distance has the nice property of being *stable*, in the sense that initially close objects will always have diagrams that are close in the bottleneck distance (Figure 1.10). For instance, in the context of two point clouds P, P' and corresponding diagrams μ, μ' built using the Čech filtration (Figure 1.8), it yields

$$\mathbf{d}_\infty(\mu, \mu') \leq d_H(P, P'),$$

where d_H stands for the Hausdorff distance between the point clouds (see Example A.21 for details).




Figure 1.10: Stability of persistence diagrams for the bottleneck distance. Here, the bottleneck distance between the two diagrams is controlled by the Hausdorff distance between the point clouds.

Limitations. Given a collection of diagrams, and thanks to the bottleneck distance, it becomes possible to compute pairwise distances between diagrams. It makes it possible to use these diagrams as the input of some simple machine learning algorithms, such as the *nearest neighbor* algorithm or the metric *multi-dimensional scaling* (MDS) one, as these algorithms only require to be able to compute distances between our observations (here, the persistence diagrams).

However, most modern machine learning algorithms ask for a *linear* structure on the space in which the data live. Unfortunately, the *space* of persistence diagrams is not linear, which essentially means that there is no simple way¹² to define the sum of two diagrams, or the multiplication of a diagram by a real number. Therefore, it is not possible to incorporate faithfully persistence diagrams in a modern machine learning pipeline, compromising their use in practical applications.

An alternative is to *embed* the space of persistence diagrams in a linear space, on which it will be possible to easily perform standard learning tasks. It means that we look for a map Φ which transforms the diagrams into *vectors*, namely an element of a linear space. Such a map is called a *vectorization*.

Various vectorizations have been proposed, with significant success in

¹²That will be compatible with the bottleneck distance

applications [Bub15, AEK⁺17, CCO17]. Nonetheless, using a vectorization raises some important questions:

- The choice of the vectorization is arbitrary. There is currently no heuristic available to choose, among all the possible vectorizations, which one will be adapted to perform well on a given learning task.
- It seems impossible¹³ to have, in the greatest generality, a *coarse* embedding, that is an embedding for which it holds, for any pair of diagrams X, Y , and two non-decreasing map $\rho_1, \rho_2 : [0, +\infty) \rightarrow [0, +\infty)$ with $\rho_1(t) \rightarrow +\infty$ when $t \rightarrow +\infty$,

$$\rho_1(\mathbf{d}_\infty(X, Y)) \leq \|\Phi(X) - \Phi(Y)\| \leq \rho_2(\mathbf{d}_\infty(X, Y)),$$

which would ensure that close diagrams have close representations, and vice-versa.

- Finally, let us recall that diagrams benefit from their interpretability. Using a vectorization, and doing the computations in the corresponding linear space, would lose this interpretability. For instance, given a map Φ and a set of diagrams $X_1 \dots X_n$, one can compute a mean in the linear space by setting $M = \frac{1}{n} \sum_{i=1}^n \Phi(X_i)$. However, it is important to note that M has, *a priori*, nothing to do with a notion of mean in the space of persistence diagrams as it may not even be in the image of Φ . This holds more generally for any computation performed in the linear space (which could lead to clustering, etc.).

Challenges. While studying vectorizations of the space of persistence diagrams remains an important research topic, we propose in this manuscript to work *in the space of persistence diagrams directly*. The goal is thus to answer the following question:

“How can we do statistics in the space of persistence diagrams?”

To answer this question, it is important to understand the geometric properties of this space. This should allow us to provide a solid theoretical framework in which statistical tools will be properly defined. Finally, these statistical tools must be usable in practice, that is we must be able to compute (or at least estimate) them.

¹³This problem is still open, but all recent results seem to point this way [BV18, CB19, BLPY19, BW20, Wag19].

1.2.4 Contributions and outline

To provide some solutions to this problem, we will rely on another field of applied mathematics: *optimal transport* theory, presented in Section 2.2. It is one of the most important contributions of this manuscript: establishing a formal connection between the geometry of the space of persistence diagrams and models used in optimal transport. This connection between these two fields is highly useful: optimal transport is a well-developed domain both from a theoretical and computational perspective and has many applications in statistical and machine learning. It turns out that most of its tools can be transposed or adapted to deal with persistence diagrams, which will allow us to obtain new theoretical results regarding the use of persistence diagrams in statistics and learning, but also to provide various algorithms and computational tools that allow us to address some difficult problems in topological data analysis, such as the estimation of Fréchet means, quantization, vectorization learning, to name a few.

In general and unless stated differently, results (theorems, propositions, etc.) for which we provide a proof are consequences of my work or joint work with my Ph.D. advisors or collaborators: Mathieu Carrière, Frédéric Chazal, Marco Cuturi, Vincent Divol, Yuichi Ike, Steve Oudot, Martin Royer, and Yuhei Umeda.

Outline and relations with the research productions during the Ph.D. thesis.

Chapter 2: Background. This chapter presents the two fields involved in this work: topological data analysis (Section 2.1) and optimal transport (Section 2.2). It aims to provide a concise introduction to both domains and a presentation of the different problems tackled in this manuscript. Note that a presentation of the algebraic construction of persistence diagrams can be found in Appendix A.

The contributions are organized into two separate parts, according to their theoretical or applied (roughly speaking, results leading to an implementation) nature.

Part I — Theory.

Chapter 3: an optimal transport framework for persistence diagrams. This chapter presents the theoretical foundations of the manuscript. It essentially corresponds to Section 3 in [DL19], under minor revisions for the *Journal of Applied and Computational Topology*. We present here how

the standard metrics used to compare persistence diagrams can be reformulated as optimal *partial* transport problems, and the new theoretical results that come out of it.

Chapter 4: A theoretical study of Fréchet means for persistence diagrams. This chapter is dedicated to the study of Fréchet means (or barycenters) of persistence diagrams from a theoretical perspective. It essentially corresponds to Section 4 in [DL19]. We prove the existence of Fréchet means in great generality, and establish a close link between those and the “Wasserstein barycenters”, their counterparts in standard optimal transport theory.

Part II — Applications.

Chapter 5: Fast algorithms for the estimation of Fréchet means of persistence diagrams. We propose an algorithm to estimate these means, which is useful especially on large scale problems. This approach has been published in the proceedings of the international conference on *Neural Information Processing Systems, 2018* [LCO18].

Chapter 6: Linear representation of persistence diagrams. This is dedicated to the study of *linear* vectorizations of the space of persistence diagrams, for which we provide an exhaustive characterization. The theoretical part (Section 6.1) comes from [DL19, §5.1]. We then provide an application in machine learning, where we introduce PERSLAY, a neural network layer¹⁴ devoted to learning optimal linear vectorizations to solve a given learning task. This work has been published in the proceedings of the international conference on *Artificial Intelligence and Statistics, 2020* [CCI+20].

Chapter 7: Complementary examples. This last chapter gathers other algorithms developed for persistence diagrams. In Section 7.1, we showcase the strength of the formalism developed in Chapter 3 in the context of random persistence diagrams. We prove convergence and stability results of the probabilistic counterpart of persistence diagrams. Section 7.2 is dedicated to the *quantization* of persistence diagrams, a useful tool that has benefits when incorporating persistence diagrams in machine learning

¹⁴including a publicly available implementation and incorporation to the [Gudhi](#) library.

pipelines. Finally, Section 7.3 proposes a simple and efficient way to estimate a shift-invariant distance between persistence diagrams.

Code and contributions to the Gudhi library. Most of the methods presented in Part II have been—or will be—incorporated into the topological data analysis library [Gudhi](#).

CHAPTER 2

BACKGROUND

Abstract

This chapter presents some background material coming from the two main research fields involved in this work: topological data analysis (Section 2.1) and optimal transport (Section 2.2). They are essentially presented in a way that serves the contents of Part I and Part II. Therefore, they do not aim at being exhaustive but rather at providing a decent state-of-the-art of the topics covered by this thesis.

2.1 Topological Data Analysis

This section is dedicated to the presentation of our objects of interest: persistence diagrams. A detailed algebraic construction of persistence diagrams—which is not required to understand the vast majority of this manuscript—can be found in Appendix A. We also refer the interested reader to [EH10, Oud15] for a thorough description. Subsection 2.1.1 below gives a concise presentation of these notions. Subsections 2.1.2 and 2.1.3 introduce various metrics between persistence diagrams and some of the properties of the resulting metric spaces.

2.1.1 Persistent homology in a nutshell

Let \mathbb{X} be a topological space, and $f: \mathbb{X} \rightarrow \mathbb{R}$ a real-valued continuous function. The t -sublevel set of (\mathbb{X}, f) is defined as

$$\mathcal{F}_t = \{x \in \mathbb{X}, f(x) \leq t\}.$$

Making t increase from $-\infty$ to $+\infty$ gives an increasing sequence of sublevel sets $(\mathcal{F}_t)_t$ called the *filtration* induced by f . To the increasing family of topological spaces $(\mathcal{F}_t)_t$ corresponds a *persistence module* $(\mathbb{V}_t)_t$, which is a family of \mathbb{K} -vector spaces (for some fixed field of coefficients \mathbb{K}) equipped with linear maps $v_s^t : \mathbb{V}_s \rightarrow \mathbb{V}_t$ for $s \leq t$ which are induced by the inclusion $\mathcal{F}_s \subset \mathcal{F}_t$ (see Appendix A for details). In particular, for $I = [b, d] \subset \overline{\mathbb{R}}$ an interval—with $\overline{\mathbb{R}} = \mathbb{R} \cup \{\pm\infty\}$ —, the *interval module* $\mathbb{I}_{[b,d]}$ is the persistence module defined by $\mathbb{I}_{[b,d]}_t = \mathbb{K}$ for $t \in I$, and $\{0\}$ otherwise, while $v_s^t = \text{id}_{\mathbb{K}}$ if $s, t \in I$ and $v_s^t = 0$ otherwise. Under mild assumptions (e.g. $\text{rk}(v_s^t)$ is finite for any s, t), a persistence module $(\mathbb{V}_t)_t$ can be *decomposed* uniquely as a direct sum of interval modules, which reads

$$\mathbb{V} = \bigoplus_{j \in J} \mathbb{I}_{[b_j, d_j]},$$

for some locally finite family of intervals $([b_j, d_j])_{j \in J}$. Therefore, \mathbb{V} is entirely described by a multiset¹ of points $\{(b, d)\}$ that belong to the extended upper half-plane

$$\overline{\mathbb{R}}_{\geq}^2 := \{(b, d) \in \overline{\mathbb{R}}^2, d \geq b\},$$

called the barcode, or *persistence diagram*, of (\mathbb{X}, f) , also denoted by $\text{Dgm}(\mathbb{X}, f)$ or simply $\text{Dgm}(f)$ if there is no ambiguity.

Intuitively, an interval $[b, d]$ appearing in $\text{Dgm}(\mathbb{X}, f)$ accounts for the presence of a topological feature (connected component, loop, cavity, etc.) that appears at scale $t = b$ in $(\mathcal{F}_t)_t$ and disappears (gets “filled”) at scale $t = d$. Equivalently, a persistence diagram can be described as a *point measure* supported on $\overline{\mathbb{R}}_{\geq}^2$, that is a measure μ of the form

$$\mu = \sum_{x \in X} n_x \delta_x,$$

where X is a locally finite subset of $\overline{\mathbb{R}}_{\geq}^2$, n_x is an integer, and δ_x denotes the Dirac mass supported on $\{x\}$. This measure-theoretic perspective on persistence diagrams is motivated in [CDSGO16] (see also Theorem A.17) and has huge benefits when studying the properties of the space of persistence diagrams, as showcased in Chapter 3.

2.1.2 Metrics between persistence diagrams

In order to use persistence diagrams in practice, we must—at the very least—be able to compare them using a metric. The standard metric used

¹Points can be repeated in the family. The number of times a point appears in a the multiset is called its *multiplicity*.




Figure 2.1: Partial matching between two diagrams. The ‘2’ indicates that the point has multiplicity 2.

to compare persistence diagrams is the bottleneck distance, defined below. It is a *partial matching* distance between the points of the diagrams, counted with multiplicity. We recall that $\overline{\mathbb{R}}_{\geq}^2$ denotes the extended closed upper half-plane. Let also $\partial\Omega$ denote the diagonal (this notation will be clarified later on, in Definition 2.9):

$$\partial\Omega := \{(t, t), t \in \mathbb{R}\}.$$

DEFINITION 2.1. Let $\mu = \sum_i \delta_{x_i}$ and $\nu = \sum_j \delta_{y_j}$ be two persistence diagrams. Let $X = (x_i)_i \subset \overline{\mathbb{R}}_{\geq}^2$ and $Y = (y_j)_j \subset \overline{\mathbb{R}}_{\geq}^2$ be the points of the diagrams, counted with multiplicity.

A partial matching γ is a subset of $(X \cup \partial\Omega) \times (Y \cup \partial\Omega)$ such that any $x_i \in X$ (resp. y_j in Y) appears exactly once as a first (resp. second) coordinate in γ (see Figure 2.1).

Let $\Gamma(\mu, \nu)$ denote the set of partial matchings between μ and ν .

The cost of a matching $\gamma \in \Gamma(\mu, \nu)$ is defined as the length of the longest edge in the matching, that is $\max\{\|x - y\|_{\infty}, (x, y) \in \gamma\}$.

The bottleneck distance between the two diagrams is then defined as the minimal cost that can be achieved by such a matching, that is

$$\mathbf{d}_{\infty}(\mu, \nu) = \inf_{\gamma \in \Gamma(\mu, \nu)} \max_{(x, y) \in \gamma} \|x - y\|_{\infty}. \quad (2.1)$$

A partial matching that realizes this infimum is said to be optimal.

REMARK 2.2. We start by noting some important remarks about these definitions.

- In the above definition, x (resp. y) belongs either to X or to the diagonal $\partial\Omega$. If $x \in X$ has finite coordinates and is matched to $\partial\Omega$, that is it belongs to a tuple $(x, y) \in \gamma$ with $y \in \partial\Omega$, then one can always assume that y is the orthogonal projection of x onto the diagonal as this would only reduce the cost.
- Similarly, although a partial matching might theoretically contain couples $(x, y) \in \partial\Omega \times \partial\Omega$, these can be removed when looking for an optimal matching as this does not increase the cost.
- Persistence diagrams might contain points that belong to $\partial\Omega$, in which case they can always be matched to the diagonal (with themselves) with a null cost. In particular, the bottleneck distance is only a pseudo-metric at this stage. This aspect will be clarified in the following subsection.
- In the standard definition of the bottleneck distance, $\|\cdot\|_\infty$ stands for the L^∞ norm in \mathbb{R}^2 , although it can be replaced with any norm $\|\cdot\|_q$ with $1 \leq q \leq \infty$. Adapting the following results is straightforward due to the equivalence of norms in \mathbb{R}^2 . When q does not play any role, we simply use the notation $\|\cdot\|$.

REMARK 2.3. Persistence diagrams may contain points with coordinates of the form $x = (b, +\infty)$ (resp. $(-\infty, d)$). The set of such points is called the essential part of the diagram. When comparing two persistence diagrams, if the cardinalities of their essential parts differ, the cost of any matching (thus the bottleneck distance) is $+\infty$.

Otherwise, the points of the respective essential parts must be matched together for the cost to be finite, and it is straightforward to observe that optimally matching those points consists in sorting the points with respect to their first (resp. second) coordinate and then using the increasing matching. It mostly means that the behavior of essential parts is obvious and independent of the rest of the diagrams.

In the rest of this manuscript, we will therefore only consider persistence diagrams with empty essential parts. Most of our theoretical and practical results can be easily adapted to take the essential parts into account separately.

Using the bottleneck distance to compare persistence diagrams is motivated by algebraic considerations detailed in Appendix A. In addition, this distance is supported by the following strong stability result, due to




Figure 2.2: Stability of the bottleneck distance for the Čech diagrams with respect to the Hausdorff distance between the point clouds.

[CSEH07, CCSG⁺09a], which states that two close filtrations induce close persistence diagrams in the bottleneck distance.

THEOREM 2.4 (*Stability theorem*). *Let $f, g : \mathbb{X} \rightarrow \mathbb{R}$ be two q -tame functions (they induce q -tame persistence modules, see Appendix A), and let $\text{Dgm}(f), \text{Dgm}(g)$ be their respective persistence diagrams. Then,*

$$\mathbf{d}_\infty(\text{Dgm}(f), \text{Dgm}(g)) \leq \|f - g\|_\infty, \quad (2.2)$$

that is the operator $f \mapsto \text{Dgm}(f)$ is 1-Lipschitz, that is:

Example A.21 in Appendix A provides a simple application in the context of diagrams built on top of point clouds. Briefly, if P and P' are two finite point clouds in \mathbb{R}^d , one has

$$\mathbf{d}_\infty(\text{Dgm}(d_P), \text{Dgm}(d_{P'})) \leq d_H(P, P'),$$

where d_P (resp. $d_{P'}$) denotes the distance to P (resp. P'), and $d_H(P, P')$ is the Hausdorff distance between P and P' : similar point clouds have similar persistence diagrams. This fact is illustrated in Figure 2.2.

Other metrics of interest. From a theoretical perspective, the bottleneck distance is a natural choice to compare persistence diagrams. However, when used for statistical and machine learning applications, it has some important drawbacks. Indeed, given two diagrams μ and ν and an optimal

partial matching γ^* , the bottleneck distance between μ and ν only depends on the longest edge in γ^* . It means that

- Any matching, as long as it does not change the length of the longest edge in γ^* , will also be optimal. There are therefore plenty of optimal matchings between two diagrams in general.
- Similarly, a slight perturbation in the points in μ , as long as it does not change the length of the longest edge of an optimal matching, will not be detected in the bottleneck distance.

In statistical and machine learning applications, discriminating information might lie in these undetected properties. Therefore, in practice, a variation of the bottleneck distance is usually preferred.

DEFINITION 2.5. *Let $1 \leq p < \infty$. The p -th distance between two diagrams μ and ν is defined as*

$$\mathbf{d}_p(\mu, \nu) = \left(\inf_{\gamma \in \Gamma(\mu, \nu)} \sum_{(x, y) \in \gamma} \|x - y\|_q^p \right)^{\frac{1}{p}}. \quad (2.3)$$

In (2.3), $\|\cdot\|_q$ stands for any q -norm, with $1 \leq q \leq \infty$. Natural choices are $q = 2$ (Euclidean ground metric), $q = p$, and $q = \infty$. This choice has barely any impact on the theoretical results presented in the manuscript. The choice $q = p$ might have benefits in numerical problems (see Chapter 5) as it leads to a *separable cost*.

While $\mathbf{d}_p \rightarrow \mathbf{d}_\infty$ as $p \rightarrow \infty$, it improves on the drawbacks mentioned above: the optimal partial matching between two diagrams (for the \mathbf{d}_p metric) is generically unique, its cost involves *all* the edges of the matching (making this distance able to account for small perturbations), etc. The p -th diagram distance also satisfies a (weaker) stability property.

DEFINITION 2.6. *A metric space \mathbb{X} is said to imply a bounded degree- q total persistence if there exists a constant $C_{\mathbb{X}}$ such that for every tame Lipschitz function $f : \mathbb{X} \rightarrow \mathbb{R}$, with $\text{Lip}(f) \leq 1$, one has $\text{Pers}_q(\text{Dgm}(f)) \leq C_{\mathbb{X}}$, where*

$$\text{Pers}_q(\text{Dgm}(f)) := d_q(\text{Dgm}(f), \mathbf{0}) = \left(\sum_{x \in \text{Dgm}(f)} \|x - \partial\Omega\|^q \right)^{\frac{1}{q}}, \quad (2.4)$$

with $\mathbf{0}$ denoting the empty diagram.

THEOREM 2.7 ([CSEHM10]). *Let \mathbb{X} be a triangulable, compact metric space that implies bounded degree- q total persistence for some $q \geq 1$, and $f, g : \mathbb{X} \rightarrow \mathbb{R}$ two tame Lipschitz functions. There exists a constant $C_{\mathbb{X}}$ that depends on \mathbb{X} such that for all $p \geq q$,*

$$\mathbf{d}_p(\text{Dgm}(f), \text{Dgm}(g)) \leq C_{\mathbb{X}}(\max \text{Lip}(f)^q, \text{Lip}(g)^q)^{\frac{1}{p}} \|f - g\|_{\infty}^{1 - \frac{q}{p}}. \quad (2.5)$$

REMARK 2.8. *The p -th distances between persistence diagrams are generally referred to as “Wasserstein distances” between persistence diagrams, a terminology that comes from optimal transport theory, see Section 2.2 below. As optimal transport will take a big place in this manuscript, we use the word “Wasserstein” to denote standard optimal transport metrics to avoid confusion.*

2.1.3 Statistics and machine learning with persistence diagrams

Let us end this section by introducing important notions for the development of this manuscript. Namely, we start by defining and stating the main properties of the *space of persistence diagrams*. We then present some statistical and machine learning problems that take place in this space.

Recall (Remark 2.3) that we only consider diagrams whose points have finite coordinates.

The space of persistence diagrams

We can now compare persistence diagrams using either the bottleneck or the p -th distances (for the sake of simplicity, we use the notation \mathbf{d}_p for both metrics). It invites us to consider the space of all possible diagrams equipped with such metrics.

First, note that the wording “metrics” is actually improper as the bottleneck and the p -th distances are only *pseudo-metrics* in the sense that one can have $\mathbf{d}_p(\mu, \nu) = 0$ while $\mu \neq \nu$. This is due to the points in the diagrams that lie on the diagonal $\partial\Omega$, as those can be matched to the diagonal with a cost of 0. This leads to a natural equivalence between diagrams: two persistence diagrams are equivalent if their off-diagonal parts coincide. The resulting equivalence classes will still be referred to as persistence diagrams. From a practical perspective, this is equivalent to simply removing the points lying on the diagonal. We can now give our core definitions.

DEFINITION 2.9. *Let us define the open upper half-plane*

$$\Omega := \{(b, d) \in \mathbb{R}^2, b < d\}, \quad (2.6)$$

and note that the diagonal $\partial\Omega := \{(t, t), t \in \mathbb{R}\}$ is its boundary.

A persistence diagram is a point measure supported on Ω , that is a measure of the form

$$\mu = \sum_{x \in X} n_x \delta_x,$$

where X is a locally finite subset of Ω , and $n_x \in \mathbb{N}$.

The set of all persistence diagrams will be denoted by \mathcal{D} .

Essentially, this definition is consistent with the algebraic one developed in the previous subsection, where diagrams are understood without their essential parts (as it is treated separately in a trivial manner) and points on the diagonal are removed (as those cannot be detected by the \mathbf{d}_p distances). These assumptions allow recovering true metric spaces on which we can provide a relevant analysis.

DEFINITION 2.10. *Let $\mu \in \mathcal{D}$, and $1 \leq p \leq \infty$. The total persistence of μ of parameter p is*

$$\text{Pers}_p(\mu) := \mathbf{d}_p(\mu, \mathbf{0}) \in [0, +\infty], \quad (2.7)$$

where $\mathbf{0}$ denotes the empty diagram.

The space of persistence diagrams of parameter p is

$$\mathcal{D}^p := \{\mu \in \mathcal{D}, \text{Pers}_p(\mu) < \infty\}. \quad (2.8)$$

This definition was introduced in the seminal papers [MMH11, TMMH14]. The assumption $\text{Pers}_p(\mu) < \infty$ in (2.8) ensures that for any $\mu, \nu \in \mathcal{D}^p$, the quantity $\mathbf{d}_p(\mu, \nu)$ is finite. It makes $(\mathcal{D}^p, \mathbf{d}_p)$ a metric space. Note the convenient expression, for $1 \leq p < \infty$,

$$\text{Pers}_p(\mu) = \sum_{x \in X} n_x \|x - \partial\Omega\|^p = \mu(\|\cdot - \partial\Omega\|^p), \quad (2.9)$$

where $\|x - \partial\Omega\|$ denotes the distance between a point $x \in \Omega$ and (its orthogonal projection onto) the diagonal $\partial\Omega$. When $p = \infty$, note that $\text{Pers}_\infty(\mu) = \sup_{x \in X} \|x - \partial\Omega\|$.

The space of persistence diagrams, at least for finite p , satisfies some fundamental properties.

PROPOSITION 2.11 ([MMH11, Theorems 6 and 12]). *Let $1 \leq p < \infty$. The space \mathcal{D}^p is a Polish metric space (i.e. is complete and separable in the metric \mathbf{d}_p).*

It implies that one can define probability distributions supported on \mathcal{D}^p , the starting point of any statistical analysis.

REMARK 2.12. *It is important to keep in mind that we allow for diagrams with possibly infinitely many points (although they must remain locally finite). This assumption is required for \mathcal{D}^p to be complete (see Example 2.13). Note that one could also recover a complete space by enforcing a uniform bound N on the number of points (mass) on the diagrams (i.e. considering the space of diagrams with at most N points). This assumption is however not suitable nor realistic in some situations. For instance,*

- *The number of points in a diagram is not a stable quantity in applications: slight modifications in the input might drastically change the number of points (by creating many points close to the diagonal, see for instance Figure 2.2). This phenomenon also occurs when sampling points (uniformly, iid) on a (unknown) shape X and looking at the Čech diagram of the sample. The Hausdorff distance between the sample and the shape X converges to 0 as the sample size $n \rightarrow \infty$, so does the diagram² thanks to the stability theorem. However the number of points in the diagram reflecting homology of dimension 1 (loops) explodes (see Figure 2.3): in this case, working with diagrams with at most N points would refrain us from producing a relevant analysis of convergence.*
- *In theory, some objects might have diagrams with infinitely many points. For instance, the (random) persistence diagram μ built from the sublevel sets of a Brownian motion has (almost surely) infinitely many points, while for any $\varepsilon > 0$, we have $\text{Pers}_{2+\varepsilon}(\mu) < \infty$ thus $\mu \in \mathcal{D}^{2+\varepsilon}$. In practice, if we want to compute the diagrams of approximated counterparts of such objects, we might expect the number of points in our diagrams to explode as the approximation error goes to 0.*

Nonetheless, the space of persistence diagrams with uniformly bounded mass has some interesting structure and can provide a relevant framework in some cases. It will be studied in detail in Section 3.2.

²for the bottleneck distance. A slightly more sophisticated analysis is required for the \mathbf{d}_p metric with $p < \infty$.




Figure 2.3: (Top) A n -sample on a Torus for $n = 100, 500, 5000$. (Bottom) The corresponding persistence diagrams in H_1 ; the title of each subplot indicates the number of points in the corresponding diagram. Although the diagram converges (in the bottleneck distance) to the diagram of the underlying torus (which has 2 points, one for each generating loop), the number of points in the diagram diverges.

EXAMPLE 2.13. Let $\mu_n = \sum_{k=1}^n \delta_{x_k}$, where $x_k = (0, 2^{-k})$. The sequence $(\mu_n)_n$ is a Cauchy sequence for the \mathbf{d}_p metric, as for $n < m$ we have $\mathbf{d}_p(\mu_n, \mu_m) \leq 2^{-(m-n)}$. However, the limit object is $\sum_{k \geq 1} \delta_{x_k}$, which has infinitely many points. See Figure 2.4 for an illustration.

REMARK 2.14. The space of persistence diagrams equipped with the bottleneck metric ($p = \infty$) is more intricate. Indeed, Eq. (2.8) does not define a separable metric space in this context (see Remark 3.25 in Chapter 3 or [BV18, Thm. 4.20]). A careful study of this space is provided in Section 3.3.

Thanks to the work [TMMH14], we know more about the geometric structure of \mathcal{D}^p in the case $p = 2$ (although their analysis easily extends to any finite p). See appendix C for a reminder of important notions in metric geometry (geodesic curves, curvature, etc.).

PROPOSITION 2.15. The space \mathcal{D}^2 is a geodesic space. It is a non-negatively curved Alexandrov space.




Figure 2.4: (Left) A Cauchy sequence of finite diagrams that converges to a diagram with infinitely many points. (Right) An example where the geodesic between two diagrams is not unique.

In Chapter 3, we will provide a new analysis of this space based on optimal transport theory. We will recover some results of [MMH11, TMMH14] using different tools and prove new properties of the space by adopting a measure-based formalism. This formalism will allow us to address some of the problems we present below.

Fréchet means of persistence diagrams

Assume we have a set of observations $X_1 \dots X_n$ from which we compute (for some filtration, etc.) a set of diagrams $\mu_1 \dots \mu_n \in \mathcal{D}^2$ (i.e. with $p = 2$). In order to summarize the topological information contained in the $(\mu_i)_i$ s, the most natural attempt is to define and compute a notion of mean, or *barycenter*. However, the space of persistence diagrams, equipped with any of the \mathbf{d}_p metrics, does not have a linear structure. Therefore, the notion of arithmetic mean used routinely in Euclidean spaces does not make sense in this space.

However, a characterization of the arithmetic mean in metric spaces is that it is the unique minimizer of $x \mapsto \sum_{i=1}^n \|x - x_i\|^2$. The latter formula only involves metrics and thus still makes sense in the context of metric spaces, in particular \mathcal{D}^2 .

DEFINITION 2.16. *Let $\mu_1 \dots \mu_n$ be n persistence diagrams. A Fréchet*

mean of $\mu_1 \dots \mu_n$ is any minimizer, should it exist, of the energy functional

$$\mathcal{E} : \nu \mapsto \frac{1}{n} \sum_{i=1}^n \mathbf{d}_2(\nu, \mu_i)^2. \quad (2.10)$$

More generally, if P is a probability distribution supported on \mathcal{D}^2 , the energy functional reads

$$\mathcal{E} : \nu \mapsto \int \mathbf{d}_2(\nu, \mu)^2 dP(\mu). \quad (2.11)$$

We commonly add the assumption $\mathcal{E}(\mathbf{0}) < \infty$, where $\mathbf{0}$ denotes the empty diagram.

The study of Fréchet means in \mathcal{D}^2 was initiated in [MMH11, TMMH14], where authors proved the existence of Fréchet means in the finite case (2.10) and in the general case (2.11) under some assumptions on the distribution P . Note that Fréchet means in \mathcal{D}^2 might not be unique, but generically are [Tur13]. From a computational perspective, [TMMH14] provided an algorithm to estimate a minimizer of (2.10), see Figure 2.5 for an illustration. However, this algorithm can return possibly bad estimates (see Chapter 5) and does not scale up on large samples.

Once again, by making use of the optimal transport formalism developed in Chapter 3, we will show how we can improve on this result. From a theoretical perspective, we will prove in Chapter 4 the existence of minimizers of (2.11) for any distribution P (satisfying a necessary finite moment assumption), extending the previous result of [MMH11]. From a practical perspective, making use of recent progress in computational optimal transport [PC17], we will propose in Chapter 5 a *convex* and *scalable* relaxation strategy to estimate minimizers of (2.10).

Persistence diagrams in random settings

A natural approach in statistics is to assume that our observations are generated, independently, by some underlying random process. In the context of topological data analysis, it leads to considering the *random* topological properties of this process.

For instance, consider the following: let ξ be a probability distribution supported on \mathbb{R}^d . Fix $n \in \mathbb{N}$. Let $\mathbb{X}_n = (X_1 \dots X_n)$ be a n -sample of law ξ (that is, $\forall i, X_i \sim \xi$, i.i.d.). One can build, using the Čech (or Rips) filtration (Example A.10) a *random persistence diagram* $\text{Dgm}(\mathbb{X}_n)$.

The topology of this sampling process is encoded by the (random) points of $\text{Dgm}(\mathbb{X}_n)$. As such, for any compact set $K \subset \Omega$, the quantity




Figure 2.5: An example of Fréchet mean of three persistence diagrams (red, blue, green) with uniform weights.

$\text{Dgm}(\mathbb{X}_n)(K)$ encodes the (random) number of points of $\text{Dgm}(\mathbb{X}_n)$ that fall into K ; it is an integer-valued random variable. A natural *deterministic* quantity to consider is then the expectation of this random variable. It defines the *expected persistence diagram* (EPD) $\mathbb{E}[\text{Dgm}(\mathbb{X}_n)]$ of the random process \mathbb{X}_n :

$$\text{For all } K \subset \Omega, \quad \mathbb{E}[\text{Dgm}(\mathbb{X}_n)](K) := \mathbb{E}[\text{Dgm}(\mathbb{X}_n)(K)]. \quad (2.12)$$

Obviously, $\mathbb{E}[\text{Dgm}(\mathbb{X}_n)]$ is not integer-valued. These objects were introduced by Chazal and Divol in [DC19], where the authors proved that under regularity condition on the sampling process ξ , the EPD is a Radon measure supported on Ω that admits a Lebesgue density supported on Ω .

How to study such objects? EPDs can be interpreted in a similar way to classical persistence diagrams: areas of high density indicate that, on average, random realizations of $\text{Dgm}(\mathbb{X}_n)$ have points in this area. However, the combinatorial definition of the diagram metrics (2.1) and (2.3) makes

them unable to handle measures with continuous support. Being able to compute distances involving expected persistence diagrams (or other measures with a continuous support) is of interest, as it would allow us to address interesting questions from a statistical viewpoint, such as

- How does the expected persistence diagram depend on the underlying law ξ generating the points?
- Given m realizations of \mathbb{X}_n , denoted by $(X_n^{(i)})_{i=1}^m$, a natural estimator of $\mathbb{E}[\text{Dgm}(\mathbb{X}_n)]$ is $\mu_n^m := \frac{1}{m} \sum_{i=1}^m \text{Dgm}(X_n^{(i)})$. Indeed, for any compact $K \subset \Omega$, $\mu_n^m(K) = \frac{1}{m} \sum_{i=1}^m \text{Dgm}(X_n^{(i)})(K) \rightarrow \mathbb{E}[\text{Dgm}(\mathbb{X}_n)(K)]$ as $m \rightarrow \infty$ by the law of large numbers. However, can we quantify how fast μ_n^m converges?

In Chapter 3, we show how the metrics \mathbf{d}_p can be extended to handle measures with continuous support. It will allow us to address the above questions in Section 7.1.

Learning with persistence diagrams

We end this section by mentioning some interesting questions that arise when one attempts to perform machine learning with persistence diagrams. First, note that as mentioned above, the lack of linear structure in the space of persistence diagrams prevents from using simple statistical descriptors faithfully, such as arithmetic means. *A fortiori*, some sophisticated machine learning pipelines, think for instance of deep neural networks, cannot be used as they generally take elements of a finite-dimensional vector space as input.

To circumvent this issue, various approaches have been proposed to embed the space of persistence diagrams into a linear space, essentially making use of *vectorizations*.

DEFINITION 2.17. *A vectorization is a map $\Phi : \mathcal{D} \rightarrow \mathcal{B}$ for some Banach space \mathcal{B} .*

A linear representation is a vectorization of the form $\Phi : \mu \mapsto \mu(f)$ for some fixed function $f : \Omega \rightarrow \mathcal{B}$, where \mathcal{B} is a Banach space.

In practice, essentially two types of vectorizations have been used:

- Through the use of kernel methods [KFH17, LY18, CCO17, RHBK15]. It consists in proposing a map $k : \mathcal{D} \times \mathcal{D} \rightarrow \mathbb{R}$ which is symmetric, positive definite. In this case, Aronszajn's theorem states that there

exists a map $\phi : \mathcal{D} \rightarrow \mathcal{H}$ to some Hilbert space \mathcal{H} (called a *reproducing kernel Hilbert space*) such that $k(\mu, \nu) = \langle \phi(\mu), \phi(\nu) \rangle_{\mathcal{H}}$. The map ϕ is implicit in general but various machine learning algorithms only make use of (expressions based on) the computable quantity $k(\cdot, \cdot)$. This approach will not be covered in detail in this manuscript.

- Through explicit mappings into finite-dimensional spaces [Bub15, AEK⁺17, BM19]. In practice, it turns out that many of these are linear representations.

The formalism we develop in this manuscript is especially adapted to study linear representations of persistence diagrams, as detailed in Chapter 6. We characterize exhaustively *continuous* linear representations, and detail in Section 6.2 how one can design a neural network layer that will learn a task-specific linear representation based on this formalism.

2.2 Optimal Transport

This section is addressed to the reader who is new to the optimal transport (OT) literature. This brief introduction will essentially focus on the important aspects of OT that will be used through the manuscript. We indicate to the interested reader three books that provide a significantly more exhaustive presentation of the field, ranked from the most theory-heavy to the more computation-oriented: [Vil08], [San15], and [PC17].

2.2.1 Monge and Kantorovich problems

We present here the optimal transport problem in its historical formulation as introduced by Gaspard Monge [Mon84], and its extension due to Kantorovich [Kan42].

Motivation and intuition

The initial motivations behind the development of the optimal transport theory were practical ones. To get the first intuition, one could consider the following problem: consider a pile of sand in some starting configuration, and suppose that you want to move it in order to reach some target configuration (for instance, to fill a hole). Of course, moving the sand is exhausting, and you may want to minimize the amount of effort put to reach your goal.




Figure 2.6: Intuitive illustration of the optimal transportation problem in 1D.

Without being completely formal yet, we can model this problem in the following way: initially, the sand is distributed in some manner, denoted by μ . The notation $d\mu(x)$ will denote the quantity (say, mass) of sand located at a given position x . Similarly, the target distribution (i.e. configuration) will be denoted by ν . *Transporting* μ onto ν , in the sense of Monge, means finding a map $T : x \mapsto T(x)$ such that for any set of locations A , one has $\nu(A) = \mu(T^{-1}(A))$, a relation which is denoted by $\nu = T_{\#}\mu$, called a *push-forward*. Note that this implicitly assumes that the distribution μ and ν must have the same total mass: at this stage, we cannot create sand *ex-nihilo*, nor make some sand disappear.

For each initial location x , going from some position x to a target position $T(x)$ has some cost, denoted by $c(x, T(x))$. For instance, c could be the distance traveled to reach x from $T(x)$, the time taken to travel, or some price. A natural assumption is that the practical cost for the user is actually proportional to the quantity of mass transported $d\mu(x)$, that is we actually pay $c(x, T(x))d\mu(x)$ when transporting a quantity of mass $d\mu(x)$ located at x onto $T(x)$. As this holds for any location x , the total cost induced by the *transport map* T when transporting μ onto ν is given by the sum of each individual cost $c(x, T(x))d\mu(x)$, yielding

$$\int_x c(x, T(x))d\mu(x).$$

Of course, one might have various ways of transporting μ onto ν , that is to pick a map T satisfying $T_{\#}\mu = \nu$, and one might pick an *optimal* one. Namely, we are looking for a map T^* that would realize

$$\inf_{T: T_{\#}\mu = \nu} \int_x c(x, T(x))d\mu(x).$$

The Monge problem

We are now ready to give the formulation of the Monge problem. First, we consider two Polish spaces \mathcal{X} and \mathcal{Y} , and two non-negative finite measures μ and ν supported on \mathcal{X} and \mathcal{Y} respectively. Let $c : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}_+$ be a lower semi-continuous non-negatively valued cost function.

The *Monge problem* between μ and ν is

$$\text{minimize}_T \left\{ \int_{x \in \mathcal{X}} c(x, T(x)) d\mu(x) : T_{\#}\mu = \nu \right\}. \quad (2.13)$$

A map T^* minimizing (2.13) is an *optimal transport map*, and its cost $\int_x c(x, T^*(x)) d\mu(x)$ is the (Monge-) *optimal transport cost* between μ and ν , denoted by $\text{OT}_c(\mu, \nu)$.

First, observe that if μ and ν do not have the same total mass, there is no transport map satisfying $T_{\#}\mu = \nu$. We will thus make this assumption in the remainder of this section. Note also that the problem is linear in the mass, that is for any $\lambda > 0$, we have $\text{OT}_c(\lambda\mu, \lambda\nu) = \lambda\text{OT}_c(\mu, \nu)$. As such, it is common (if not systematic) in the optimal transport literature to consider that μ, ν are *probability measures* on their respective spaces, that is non-negative Radon measures of mass 1.

Unfortunately, in many situations, looking for transport maps from μ to ν is too restrictive. Indeed, this formalism enforces that the *whole* mass located at $x \in \mathcal{X}$ is transported to a *single* target location $T(x) \in \mathcal{Y}$: splitting the mass is not allowed. As such, there are simple situations where there is no transport map. For instance, fix $\mathcal{X} = \mathcal{Y} = \mathbb{R}$ and consider $\mu = \delta_0$ the Dirac mass located at 0, and $\nu = \frac{1}{2}(\delta_{-1} + \delta_1)$. In this context, there is no T satisfying $T_{\#}\mu = \nu$; see also Figure 2.7 for an illustration. Actually, even when the set of transport maps is not empty, it might happen that no transport map is optimal, see [Vil08, Example 4.9] for instance.

To overcome this limitation, one resorts to the *Kantorovich* formulation of the transport problem. Essentially, this extension of the Monge problem consists in optimizing on *transport plans* instead of transport maps. A transport plan is a measure π supported on the product space $\mathcal{X} \times \mathcal{Y}$ whose natural interpretation is the following: $d\pi(x, y)$ represents the fraction of mass initially located at x that is transported to y . In addition, we ask the transport plan to actually transport μ onto ν , which is encoded by the *marginal constraints* that read, for all Borel subsets $A \subset \mathcal{X}, B \subset \mathcal{Y}$,

$$\mu(A) = \int_{x \in A} \int_{y \in \mathcal{Y}} d\pi(x, y), \quad \nu(B) = \int_{y \in B} \int_{x \in \mathcal{X}} d\pi(x, y). \quad (2.14)$$

The set of transport plans between μ and ν is called the *transportation polytope* between μ and ν and will be denoted by $\Pi(\mu, \nu)$. The Kantorovich




Figure 2.7: (Left) A basic example where there exists no transport map between μ and ν . (Right) An example where there exists an optimal Monge map (dashed), but performs poorly compared to an optimal transport plan (that would be able to split the $(1/2)$ red mass and transport it onto the $(1/4)$ blue masses).

formulation of the optimal transport problem finally reads

$$\text{minimize}_{\pi \in \Pi(\mu, \nu)} \iint_{\mathcal{X} \times \mathcal{Y}} c(x, y) d\pi(x, y). \quad (2.15)$$

As in the Monge transportation problem, a minimizer of (2.15) is an *optimal transport plan*, and the infimum value is the (Kantorovich-) *optimal transport cost* between μ and ν for the cost c , also denoted by $\text{OT}_c(\mu, \nu)$.

REMARK 2.18. *Unlike the Monge problem, optimal transport plans for the Kantorovich problem always exist [Vil08, Theorem 4.1]. In general, it might happen that an optimal transport map exists but achieves a much larger cost than an optimal transport plan (see Figure 2.7). However, if μ is atomless, although an optimal transport map might not exist, the infimum value in the Monge problems coincide with the minimum value of the Kantorovich problems [Pra07, Theorem B].*

Finally, let us note two important situations where optimal transport maps and optimal transport plans do coincide.

PROPOSITION 2.19 ([Vil08, Theorem 9.4, Theorem 10.41]). *If $\mathcal{X} = \mathcal{Y} = \mathbb{R}^d$ and $c(x, y) = \|x - y\|^2$ (the so-called quadratic cost), and if μ admits a density with respect to the Lebesgue measure, then there exists a unique optimal transport plan, which is actually an optimal transport map, which eventually turns out to be the (sub)gradient of a convex function.*

PROPOSITION 2.20 (Birkhoff-von Neumann). *If μ and ν are discrete measures with uniform weights, that is of the form $\mu = \sum_{i=1}^n \delta_{x_i}$ and $\nu = \sum_{j=1}^n \delta_{y_j}$ for some integer n , then there exists an optimal transport*

plan between μ and ν which is a transport map; that is a bijection between the supports of μ and ν , also known as an optimal matching.

2.2.2 Wasserstein distances and Wasserstein spaces

A natural context where one can apply optimal transport is the following: assume that the (Polish) space \mathcal{X} is equipped with a metric d , and consider some $1 \leq p < \infty$. The p -Wasserstein distance between μ and ν is defined as the optimal transport cost (to the power $1/p$) of transporting μ onto ν for the cost $c(x, y) = d(x, y)^p$, that is

$$W_{p,d}(\mu, \nu) := \left(\inf_{\pi \in \Pi(\mu, \nu)} \iint_{\mathcal{X} \times \mathcal{X}} d(x, y)^p d\pi(x, y) \right)^{\frac{1}{p}}, \quad (2.16)$$

When there is no ambiguity on the distance d used, we simply write W_p instead of $W_{p,d}$. In order to have W_p finite, μ and ν are required to have a finite p -th moment, that is there exists $x_0 \in \mathcal{X}$ such that $\int_{\mathcal{X}} d(x, x_0)^p d\mu(x)$ (resp. $d\nu$) is finite. A nice observation is that W_p is a metric over the set of such probability measures, referred to as $\mathcal{W}^p(\mathcal{X})$ and called the p -Wasserstein space over (\mathcal{X}, d) . A first interesting remark that can be done is the simple equality

$$W_{p,d}(\delta_x, \delta_y) = d(x, y),$$

for any $x, y \in \mathcal{X}$. Thus, the Wasserstein space $\mathcal{W}^p(\mathcal{X})$ contains an isometric copy of the ground space \mathcal{X} by simply looking at Dirac masses. From a theoretical viewpoint, the Wasserstein space over (\mathcal{X}, d) is of interest as many geometric properties of (\mathcal{X}, d) are reflected in $\mathcal{W}^p(\mathcal{X})$. For instance, if (\mathcal{X}, d) is a Polish [Vil08, Theorem 6.18], or a compact [Vil08, Remark 6.19], or a geodesic [Vil08, Corollary 7.22] space, so is $\mathcal{W}^p(\mathcal{X})$. Therefore, when dealing with a collection of probability measures $\mu_1 \dots \mu_N$ supported on some “nice” ground space (\mathcal{X}, d) , think for instance of \mathbb{R}^d equipped with the Euclidean metric, one can build statistical summaries of the $(\mu_i)_i$ s by working in the Wasserstein space. Such statistics are likely to reflect the geometry of the ground space in some sense (see Example 2.21).

EXAMPLE 2.21. *Let $\mu_1 \dots \mu_N$ be a collection of probability measures supported on \mathbb{R}^d with finite second moment. A Wasserstein barycenter of $\mu_1 \dots \mu_N$ (also called a Fréchet mean) is any minimizer (in $\mathcal{W}^2(\mathbb{R}^d)$) of the energy functional*

$$\mathcal{E} : \mu \mapsto \sum_{i=1}^N \lambda_i W_2^2(\mu, \mu_i), \quad (2.17)$$

where the λ_i s are non-negative weights that sum to 1. This problem was studied in the seminal work [AC11] where authors proved the existence of solutions. From a theoretical viewpoint, this result was extended in [LGL16] to the case of averaging infinitely many probability measures on $\mathcal{W}^p(\mathbb{R}^d)$. Namely, given a probability distribution P supported on $\mathcal{W}^p(\mathbb{R}^d)$ (a distribution of probability measures), they proved the existence of solutions to

$$\mathcal{E} : \mu \mapsto \int W_2^2(\mu, \nu) dP(\nu), \quad (2.18)$$

and they also proved that minimizers of (2.17) must converge to minimizers of (2.18) as $N \rightarrow \infty$ —under the assumption that the $\mu_i \sim P$, i.i.d.

From a computational perspective, based on entropic regularization (see Section 2.2.4 below) of the Wasserstein metrics, efficient methods to estimate solutions of (2.17) have been proposed in a series of works in the mid 2010’s [CD14, SDGP⁺15, CP16], see Figure 2.8 for an illustration. Wasserstein barycenters are still actively studied; among interesting research branches, we point out the work [GPRS19] where authors study the convergence rates of minimizers of (2.17) to minimizers of (2.18) as $N \rightarrow \infty$ (where $\mu_i \sim P$).

2.2.3 Unbalanced optimal transport

Wasserstein distances and \mathbf{d}_p metrics between persistence diagrams defined in Eq. (2.3) share the key idea of defining a distance by minimizing a cost over some matchings. However, the set of transport plans $\Pi(\mu, \nu)$ between two measures is non-empty if and only if the two measures have the same mass, while persistence diagrams with different masses (number of points) can be compared, making a crucial difference between the W_p and \mathbf{d}_p metrics.

Extending optimal transport to measures of different masses, generally referred to as *unbalanced optimal transport*, has been addressed by different authors [Fig10, CPSV18, KMV⁺16], and we refer to the PhD thesis manuscript of L.Chizat [Chi17] for a general presentation. We present in the following two approaches to deal with unbalanced measures that will be of interest to handle persistence diagrams.

Regularization of marginal constraints

Consider a Polish space \mathcal{X} and two non-negative finite Radon measures μ, ν supported on \mathcal{X} . As mentioned above, the transportation polytope $\Pi(\mu, \nu)$ is empty if (and only if) μ and ν do not have the same masses. Intuitively,




Figure 2.8: Wasserstein barycenters of three measures for different weights.

it means that in this context, satisfying the marginal constraints (2.14) is too strong of a requirement. A natural workaround is then to relax this requirement: instead of asking the marginals of a measure π on $\mathcal{X} \times \mathcal{X}$ to be *exactly* μ and ν , we penalize the distance between the marginals and the “targets” μ and ν . Formally, this is done using a *Csiszár-divergence* (or *φ -divergence*).

Let $\varphi : \mathbb{R} \rightarrow \mathbb{R}_+$ be an *entropy function*, that is a convex lower semi-continuous function, satisfying $\varphi(p) = +\infty$ if $p < 0$ and $\varphi(1) = 0$. Define $\varphi'_\infty = \lim_{p \rightarrow \infty} \varphi(p)/p$. The divergence associated to φ between two non-negative Radon measures is defined as

$$D_\varphi(\alpha|\beta) := \int_{\mathcal{X}} \varphi\left(\frac{d\alpha}{d\beta}\right) d\beta(x) + \varphi'_\infty \int_{\mathcal{X}} d\alpha', \quad (2.19)$$

where $\frac{d\alpha}{d\beta}$ is the *Radon-Nikodym* decomposition of α with respect to β (see Appendix B), with the standard convention $\infty \times 0 = 0$ if $\varphi'_\infty = \infty$. The

unbalanced transport problem between μ and ν is then defined as

$$\underset{\pi \in \mathcal{M}(\mathcal{X} \times \mathcal{X})}{\text{minimize}} \iint c(x, y) d\pi + \rho D_\varphi(\pi^{(1)} | \mu) + \rho D_\varphi(\pi^{(2)} | \nu), \quad (2.20)$$

where $\pi^{(1)}, \pi^{(2)}$ denote the first and second marginals of π respectively, and $\rho > 0$ is a parameter controlling the strength of the (regularized) marginal constraints.

EXAMPLE 2.22. *Some choices of φ allow recovering well-known notions of divergence between measures. For instance,*

- $\varphi : z \mapsto z \log(z) - z + 1$ defines $D_\varphi = \text{KL}$, the Kullback-Leibler divergence.
- $\varphi : z \mapsto |1 - z|$ defines the total variation between measures.

Note also that taking $\varphi : z \mapsto 0$ if $z = 1$, and $\varphi(z) = +\infty$ otherwise, define the convex indicator function of 1, in which case we recover the balanced OT problem (2.15) as a particular case of (2.20).

This formalism can be appealing when attempting to provide an optimal transport formulation of metrics between persistence diagrams. However, it is important to observe that the penalization $D_\varphi(\pi^{(1)} | \mu)$ is *independent of the geometry of the ground space*. When comparing persistence diagrams (2.3), in contrast, we considered that unmatched points were transported to their orthogonal projections onto the diagonal; so the cost of not transporting a point depends on its location.

Optimal transport with boundary

An alternative approach, that is not encompassed by (2.20), was proposed by Figalli and Gigli in [FG10]. In this work, authors propose to extend Wasserstein distances to Radon measures supported on a bounded open proper subset \mathcal{X} of \mathbb{R}^d , whose boundary is denoted by $\partial\mathcal{X}$ (and $\bar{\mathcal{X}} := \mathcal{X} \sqcup \partial\mathcal{X}$).

Informally, they propose to use the boundary $\partial\mathcal{X}$ as an infinite reservoir of mass. To do that, they relax the marginal constraints, proposing to match an element of mass $d\mu(x)$ either to some $d\nu(y)$, with cost $d(x, y)^p$, or to the boundary of the space $\partial\mathcal{X}$, with cost $d(x, \partial\mathcal{X}) := \min_{x' \in \partial\mathcal{X}} d(x, x')^p$ (and similarly for ν). This is formalized by the following definition.

DEFINITION 2.23. [FG10, Problem 1.1] Let $p \in [1, +\infty)$. Let μ, ν be two Radon measures supported on \mathcal{X} satisfying

$$\int_{\mathcal{X}} d(x, \partial\mathcal{X})^p d\mu(x) < +\infty, \quad \int_{\mathcal{X}} d(x, \partial\mathcal{X})^p d\nu(x) < +\infty. \quad (2.21)$$

The set of admissible transport plans (or couplings) $\text{Adm}(\mu, \nu)$ is defined as the set of measures π on $\overline{\mathcal{X}} \times \overline{\mathcal{X}}$ satisfying for all Borel sets $A, B \subset \mathcal{X}$,

$$\pi(A \times \overline{\mathcal{X}}) = \mu(A) \quad \text{and} \quad \pi(\overline{\mathcal{X}} \times B) = \nu(B).$$

The cost of $\pi \in \text{Adm}(\mu, \nu)$ is defined as

$$C_p(\pi) := \iint_{\overline{\mathcal{X}} \times \overline{\mathcal{X}}} d(x, y)^p d\pi(x, y). \quad (2.22)$$

The Optimal Transport (with boundary) distance $\text{OT}_p(\mu, \nu)$ is then defined as

$$\text{OT}_p(\mu, \nu) := \left(\inf_{\pi \in \text{Adm}(\mu, \nu)} C_p(\pi) \right)^{1/p}. \quad (2.23)$$

As in the balanced OT formalism (2.15), plans $\pi \in \text{Adm}(\mu, \nu)$ realizing the infimum in (2.23) are called optimal. The set of optimal transport plans between μ and ν for the cost $(x, y) \mapsto d(x, y)^p$ is denoted by $\text{Opt}_p(\mu, \nu)$.

As we will see later in Chapter 3, this formulation is exactly the one we need to bridge the gap between persistence diagram metrics and optimal transport formulations. At this point, we simply put the emphasis on some important facts:

- OT_p is well-defined even between Radon measures of infinite mass, as long as they satisfy (2.21). This property can be seen as an analog to the finite moment condition involved in the definition of the Wasserstein distance in Section 2.2.2.
- In their work, Figalli and Gigli supposed the ground space \mathcal{X} to be bounded. Of course, the upper half-plane Ω on which persistence diagrams are supported is not bounded. This does not prevent us from defining OT_p in this context but will require additional care in the proofs.

We end this section by introducing a useful notion that will be used in Chapter 3, which shows how to build an element of $\text{Adm}(\mu, \nu)$ given a map $f : \overline{\mathcal{X}} \rightarrow \overline{\mathcal{X}}$ satisfying some balance condition (see Figure 2.9).




Figure 2.9: A transport map f must satisfy that the mass $\nu(B)$ (light blue) is the sum of the mass $\mu(f^{-1}(B) \cap \mathcal{X})$ given by μ that is transported by f onto B (light red) and the mass $\nu(B \cap f(\partial\mathcal{X}))$ coming from $\partial\mathcal{X}$ and transported by f onto B .

DEFINITION 2.24. Let $\mu, \nu \in \mathcal{M}$. Consider $f : \overline{\mathcal{X}} \rightarrow \overline{\mathcal{X}}$ a measurable function satisfying for all Borel set $B \subset \mathcal{X}$

$$\mu(f^{-1}(B) \cap \mathcal{X}) + \nu(B \cap f(\partial\mathcal{X})) = \nu(B). \quad (2.24)$$

Define for all Borel sets $A, B \subset \overline{\mathcal{X}}$,

$$\pi(A \times B) = \mu(f^{-1}(B) \cap \mathcal{X} \cap A) + \nu(\mathcal{X} \cap B \cap f(A \cap \partial\mathcal{X})). \quad (2.25)$$

π is called the transport plan induced by the transport map f .

One can easily check that we have indeed $\pi(A \times \overline{\mathcal{X}}) = \mu(A)$ and $\pi(\overline{\mathcal{X}} \times B) = \nu(B)$ for any Borel sets $A, B \subset \mathcal{X}$, so that $\pi \in \text{Adm}(\mu, \nu)$ (see Figure 2.9).

2.2.4 Entropic regularization of optimal transport

The last variation of the optimal transport problem we need to introduce in this background section is the *entropic regularization* of optimal transport, an idea that can be traced back to the Schrödinger problem [Sch31, Léo14]. This approach was highly popularized by Cuturi [Cut13], showing that the (unique!) solution of the regularized problem introduced below can be approximated at a lightspeed rate.

Definition

Consider μ, ν , two probability measures supported on some space \mathcal{X} equipped with some cost function c (typically $c(x, y) = d(x, y)^p$, for $p \geq 1$). The transportation polytope $\Pi(\mu, \nu)$ is not empty: it contains at least the *product measure* $\mu \otimes \nu$, defined as $\mu \otimes \nu(A, B) := \mu(A)\nu(B)$ for all Borel $A, B \subset \mathcal{X}$. Interestingly, it turns out that $\mu \otimes \nu$ is the element of $\Pi(\mu, \nu)$ that has minimal *negentropy* h , defined as

$$h : \pi \mapsto \iint (\log(\pi) - 1) d\pi,$$

with the standard convention $0 \times \log(0) = 0$. As such, an intuitive understanding one can have about the entropic regularization of optimal transport is the following. On the one hand, there is (at least) one optimal transport plan π^* , but this might be “hard” to find. On the other hand, there is the trivial transport plan $\mu \otimes \nu$, which is straightforward to compute, but which will be suboptimal in general. Why not interpolate between both? It turns out that $\mu \otimes \nu$ is the element of $\Pi(\mu, \nu)$ that has maximal entropy (minimal negentropy). It invites us to consider, for some regularization parameter $\varepsilon > 0$,

$$W_c^\varepsilon(\mu, \nu) := \inf_{\pi \in \Pi(\mu, \nu)} \iint c(x, y) d\pi + \varepsilon h(\pi). \quad (2.26)$$

A slight variation of this definition, used in recent works in optimal transport [RTC17, GPC18, FSV⁺19, SFV⁺19], consists in changing the negentropy h for the *Kullback-Leibler divergence* to the product measure $\mu \otimes \nu$, reading

$$\text{OT}_c^\varepsilon(\mu, \nu) := \inf_{\pi \in \Pi(\mu, \nu)} \iint c(x, y) d\pi + \varepsilon \text{KL}(\pi | \mu \otimes \nu), \quad (2.27)$$

where we recall that, as μ, ν have mass 1,

$$\text{KL}(\pi | \mu \otimes \nu) = \iint \log \left(\frac{d\pi}{d\mu \otimes \nu} \right) d\pi.$$

In both cases, the first consequence of introducing the regularization term is that it makes the minimization problem *strictly* convex, ensuring that there is a *unique* solution π^ε . Furthermore, $\pi^\varepsilon \rightarrow \pi^*$ as $\varepsilon \rightarrow 0$, where π^* is the optimal transport plan between μ and ν (for the non-regularized problem) that has maximal entropy.

REMARK 2.25. As explained in [PC17, Remark 4.2], using the Kullback-Leibler divergence instead of the negentropy is nothing but a change of “reference measure”. Indeed, the negentropy h is (up to a constant factor) the Kullback-Leibler divergence with respect to the Lebesgue measure. Using $\mu \otimes \nu$ as a reference measure is motivated by the following fact: as the regularization parameter ε goes to ∞ , the unique solution π^ε (either of (2.26) or (2.27)) is expected to get closer to $\mu \otimes \nu$ (as it is the measure in $\Pi(\mu, \nu)$ of largest entropy). Using any other reference measure would make the regularized transport cost to diverge as $\varepsilon \rightarrow \infty$. In contrast, one has

$$\text{OT}_c^\varepsilon(\mu, \nu) \rightarrow \iint c(x, y) d\mu(x) d\nu(y)$$

as $\varepsilon \rightarrow \infty$. A proof of this fact can be found in the Ph.D. thesis of Aude Genevay (see [Gen19, Theorem 10]).

Eventually, it turns out that the map $(\mu, \nu) \mapsto \pi^\varepsilon$ is smooth. This smoothness is at the core of some variational algorithms to solve optimization problems in the Wasserstein space, such as computing Wasserstein barycenters. We will present in Chapter 5 an application of this property in the context of persistence diagrams.

Computational aspect

Nonetheless, what makes the entropic regularization so popular in computational optimal transport nowadays lies in its computational strengths. For the sake of exposition, and as we target numerical applications here, let us expose these strengths in the *finite-dimensional setting*, that is when the measures of interest have finite support. Namely, let $\mu = \sum_{i=1}^n a_i \delta_{x_i}$ and $\nu = \sum_{j=1}^m b_j \delta_{y_j}$, where the $a = (a_i)_i \in \mathbb{R}_+^n$ and $b = (b_j)_j \in \mathbb{R}_+^m$ are non-negative weights that sum to 1, and $(x_i)_i, (y_j)_j$ are points in \mathbb{R}^d . Finally, let $C \in \mathbb{R}^{n \times m}$ denote the matrix $(c(x_i, y_j))_{ij}$ (typically $c(x_i, y_j) = \|x_i - y_j\|_2^2$, but any choice would work in the following).

In this context, (2.26) reads³

$$W_c^\varepsilon(\mu, \nu) = \inf_{P \in \mathbb{R}^{n \times m}} \left\{ \langle P, C \rangle + \varepsilon \sum_{ij} P_{ij} (\log(P_{ij}) - 1), \right. \\ \left. P \mathbf{1}_m = a, P^T \mathbf{1}_n = b \right\}, \quad (2.28)$$

where $\langle P, C \rangle = \sum_{ij} P_{ij} C_{ij}$ denotes the Frobenius dot product between P and C , and $\mathbf{1}_k$ denotes the vector $(1 \dots 1) \in \mathbb{R}^k$. Using standard duality

³Similar computational results exist for the variation OT_c^ε , not detailed here for the sake of concision. See [FSV⁺19] for details.

results, this problem is equivalent to the following unconstrained maximization problem

$$W_c^\varepsilon(\mu, \nu) = \sup_{f \in \mathbb{R}^n, g \in \mathbb{R}^m} \langle a, f \rangle + \langle b, g \rangle - \varepsilon \sum_{ij} \exp\left(\frac{f_i + g_j - C_{ij}}{\varepsilon}\right), \quad (2.29)$$

with the primal-dual relation between a minimizer P^* in (2.28) and maximizers (f^*, g^*) of (2.29):

$$P^* = \exp\left(\frac{f_i + g_j - C_{ij}}{\varepsilon}\right). \quad (2.30)$$

Introducing the variables $u = \exp(f/\varepsilon)$, $v = \exp(g/\varepsilon)$ and $K = \exp(-C/\varepsilon)$ (where exponentiation must be understood term-wise), the first order conditions read (term-wise division)

$$\begin{cases} u = \frac{a}{Kv} \\ v = \frac{b}{K^T u} \end{cases}. \quad (2.31)$$

Said differently, an optimal pair (u, v) is a fixed point of the *Sinkhorn map*

$$S : (u, v) \mapsto \left(\frac{a}{Kv}, \frac{b}{K^T u}\right). \quad (2.32)$$

Furthermore, the sequence $(u_t, v_t)_t$ defined for arbitrary (u_0, v_0) as

$$\begin{cases} u_{t+1} \leftarrow \frac{a}{Kv_t} \\ v_{t+1} \leftarrow \frac{b}{K^T u_{t+1}} \end{cases} \quad (2.33)$$

converges toward a such a fixed point (u, v) . This gives a fairly simple algorithm to approximate a solution of (2.28) (both in terms of minimum and minimizer). The iterations involved in (2.33) only involve basic matrix manipulations and are thus especially suited to be run on modern hardware such as a GPU.

This algorithm can have even more computational benefits when dealing with *histograms* (think of images for instance), that is measures supported on a common finite set of cardinality n . A histogram is thus a vector \mathbb{R}_+^n whose coordinates sum to one. Assume that one wants to solve N instances of (2.28). It involves N pairs of variables $(u^{(1)}, v^{(1)}) \dots (u^{(N)}, v^{(N)})$ but a *single* cost matrix $C \in \mathbb{R}^{n \times n}$. As such, one can run the iterations (2.33) only once, using $u^{(1\dots N)} := (u^{(1)} \dots u^{(N)}) \in \mathbb{R}^{n \times N}$ and $v^{(1\dots N)} = (v^{(1)} \dots v^{(N)})$. This makes the Sinkhorn algorithm especially efficient when one wants to estimate a large number of transport distances between histograms.

Sinkhorn divergences

Although regularized optimal transport was initially popularized for its computational properties, some effort has been recently dedicated to providing a better understanding of its theoretical behavior. In [GPC18], Genevay et al. introduced the notion of *Sinkhorn divergence*, namely they define:

$$S_c^\varepsilon(\mu, \nu) = \text{OT}_c^\varepsilon(\mu, \nu) - \frac{1}{2}\text{OT}_c^\varepsilon(\mu, \mu) - \frac{1}{2}\text{OT}_c^\varepsilon(\nu, \nu). \quad (2.34)$$

The initial idea behind this definition is that $\text{OT}_c^\varepsilon(\mu, \mu)$ does not necessarily equal 0, while $S_c^\varepsilon(\mu, \mu)$ obviously does. Furthermore, while it was well-known that $\text{OT}_c^\varepsilon(\mu, \nu) \rightarrow \text{OT}_c(\mu, \nu)$ as $\varepsilon \rightarrow 0$ (the same hold with W_c^ε or S_c^ε), authors in [RTC17] and [GPC18] observe that

$$\begin{aligned} \lim_{\varepsilon \rightarrow +\infty} S_c^\varepsilon(\mu, \nu) &= - \iint c(x, y) d(\mu - \nu)(x) d(\mu - \nu)(y) \\ &=: \text{MMD}(\mu, \nu), \end{aligned}$$

where MMD stands for the *maximum mean discrepancy* [GBR⁺12], another notion of distance between probability measures, making the Sinkhorn divergence an interpolation between two common ways to compare probability measures. Feydy et al. [FSV⁺19, Theorem 1] strengthened this fact by showing, under some assumptions on c and for measures supported on a common compact domain, that although S_c^ε is not a distance (it does not satisfy the triangle inequality), one has for all probability measures $(\mu_n)_n, \mu, \nu$,

$$\begin{aligned} S_c^\varepsilon(\mu, \nu) &\geq S_c^\varepsilon(\mu, \mu) \geq 0, \\ S_c^\varepsilon(\mu, \nu) &= 0 \Leftrightarrow \mu = \nu, \\ S_c^\varepsilon(\mu_n, \mu) &\rightarrow 0 \Leftrightarrow \text{OT}_c(\mu_n, \mu) \rightarrow 0. \end{aligned}$$

Thus, the Sinkhorn divergence induces a good notion of *divergence* between probability measures. Finally, let us mention that Sinkhorn divergences have been extended to the unbalanced case (Section 2.2.3) in the recent work of Séjourné et al. [SFV⁺19].

2.3 Notations

We will use the following notations throughout the manuscript.

Let Ω denote the open half-plane

$$\Omega := \{(b, d) \in \mathbb{R}^2, b < d\}, \quad (2.35)$$

and let $\partial\Omega$ denote its boundary, namely the diagonal

$$\partial\Omega := \{(t, t), t \in \mathbb{R}\}. \quad (2.36)$$

For $x \in \Omega$, $\mathbf{p}_{\partial\Omega}(x)$ denotes its orthogonal projection onto the diagonal $\partial\Omega$. Let $\bar{\Omega}$ denote $\Omega \sqcup \partial\Omega$. The function $d : \bar{\Omega} \times \bar{\Omega} \rightarrow \mathbb{R}$ denotes the *ground metric* on $\bar{\Omega}$, typically $d(x, y) := \|x - y\|_q$ for some $q \in [1, +\infty]$. We will also use the notation

$$E_\Omega := \bar{\Omega} \times \bar{\Omega} \setminus \partial\Omega \times \partial\Omega.$$

The *space of persistence measures* of parameter $1 \leq p \leq \infty$ is the space of non-negative Radon measures supported on Ω that have finite *total persistence- p* , that is

$$\mathcal{M}^p := \left\{ \mu, \text{Pers}_p(\mu) := \int_\Omega d(x, \partial\Omega)^p d\mu(x) < \infty \right\}, \quad (2.37)$$

where in the case $p = \infty$ the total persistence reads

$$\text{Pers}_\infty(\mu) = \sup_{x \in \text{spt}(\mu)} d(x, \partial\Omega),$$

also known as μ -essential supremum of the function $x \mapsto d(x, \partial\Omega)$. It is equipped with an optimal partial transport metric denoted by OT_p .

The *space of persistence diagrams* of parameter p is the subset $\mathcal{D}^p \subset \mathcal{M}^p$ consisting only of *point measures*, that is measures μ of the form

$$\mu = \sum_{x \in X} n_x \delta_x \quad (2.38)$$

where $X \subset \Omega$ is a locally finite set, and $n_x \in \mathbb{N}$. It is equipped with a partial matching metric \mathbf{d}_p , although we show that $\mathbf{d}_p = \text{OT}_p$ over \mathcal{D}^p .

Given a Polish metric space (\mathcal{X}, d) , and a parameter $1 \leq p \leq \infty$, let $W_{p,d}$ (or W_p if there is no ambiguity on the ground metric d) denote the p -Wasserstein distance between probability measures (of finite p -th moment) supported on \mathcal{X} ; and $\mathcal{W}^p(\mathcal{X}, d)$ or simply $\mathcal{W}^p(\mathcal{X})$ the space of such probability measures equipped with the metric $W_{p,d}$.

Part I
Theory

CHAPTER 3

PERSISTENCE DIAGRAMS AND MEASURES, AN OPTIMAL TRANSPORT VIEWPOINT

Abstract

This chapter is dedicated to a theoretical study of spaces of persistence diagrams by adopting an optimal transport viewpoint. By establishing a strong connection between diagram metrics and optimal transport ones, we obtain new results on these spaces, in particular, a very useful characterization of convergence of persistence diagrams in terms of measure theory. Furthermore, this formalism allows us to define an extension of the diagram metrics to any measure supported on the open half-plane, not only those with discrete support such as diagrams, offering the possibility to deal with more complicated objects such as random persistence diagrams (see Section 2.1.3). The strength of this formalism will be showcased through various applications in Part II.

3.1 General properties of \mathcal{M}^p

Let Ω denote the open half-plane

$$\Omega := \{(b, d) \in \mathbb{R}^2, b > d\}, \quad (3.1)$$

and let $\partial\Omega$ denote its boundary, namely the diagonal

$$\partial\Omega := \{(t, t), t \in \mathbb{R}\}. \quad (3.2)$$

Eventually, let $\bar{\Omega}$ denote $\Omega \sqcup \partial\Omega$.

For now, we fix a parameter $p \in [1, +\infty)$. The particular case $p = \infty$ is studied in Section 3.3. Recall from Definition 2.9 in Section 2.1 that persistence diagrams are defined as locally finite point measures supported on Ω , that is of the form $\sum_{x \in X} n_x \delta_x$, where $X \subset \Omega$ is locally finite, and $n_x \in \mathbb{N}$ for all $x \in X$. The space of persistence diagrams is equipped with metrics \mathbf{d}_p , defining a metric space $(\mathcal{D}^p, \mathbf{d}_p)$. A diagram $\mu \in \mathcal{D}^p$ might have infinitely many points, but we request that $\text{Pers}_p(\mu) := \sum_{x \in X} n_x d(x, \partial\Omega)^p < \infty$, where $d(\cdot, \cdot)$ denotes the ground distance on \mathbb{R}^2 , typically $\|\cdot - \cdot\|_q$ for $q \in [1, +\infty]$ (the choice of q does not play any important role). This condition is equivalent to saying that μ must be at a finite distance from the empty diagram, i.e. $\mathbf{d}_p(\mu, \mathbf{0}) < \infty$.

Considering diagrams as measures invites us to consider more general measures supported on Ω . Let thus $\mathcal{M}(\Omega)$ denote the set of non-negative Radon measures supported on Ω (see Appendix B for details in measure theory). Without further ado, we introduce the space of *persistence measures*

$$\mathcal{M}^p := \left\{ \mu \in \mathcal{M}(\Omega), \text{Pers}_p(\mu) := \int_x d(x, \partial\Omega)^p d\mu(x) < \infty \right\}. \quad (3.3)$$

We equip \mathcal{M}^p with the metric OT_p (see Section 2.2.3 for a more detailed presentation) defined as

$$\text{OT}_p(\mu, \nu) = \inf_{\pi \in \text{Adm}(\mu, \nu)} \iint_{\overline{\Omega} \times \overline{\Omega}} d(x, y)^p d\pi(x, y), \quad (3.4)$$

where $\text{Adm}(\mu, \nu)$ is the set of *admissible plans* between μ and ν , that is the set of measures supported on $\overline{\Omega} \times \overline{\Omega}$ satisfying:

$$\text{for all Borel sets } A, B \subset \Omega, \pi(A, \overline{\Omega}) = \mu(A), \pi(\overline{\Omega}, B) = \nu(B). \quad (3.5)$$

REMARK 3.1. *These metrics have already been introduced by Figalli and Gigli in [FG10]. In their work, authors however consider the ground space (here, the open half-plane Ω) to be compact (which is obviously not the case here). Some of the results presented below have counterparts in the work of Figalli and Gigli. Sometimes, the corresponding proofs could be immediately adapted to our framework, sometimes it required to develop new techniques (and, of course, some results just do not hold on non-compact domains).*

For the ease of readability, we have delayed some proofs to a separate section (§3.5). We explicitly mention for each result if its proof straightforwardly adapts from one of [FG10] or if it can be considered as a new result or a new proof.

PROPOSITION 3.2. *Let $\mu, \nu \in \mathcal{M}^p$. The set of transport plans $\text{Adm}(\mu, \nu)$ is sequentially compact for the vague topology on $E_\Omega := \overline{\Omega} \times \overline{\Omega} \setminus \partial\Omega \times \partial\Omega$. Moreover, for this topology,*

- $\pi \in \text{Adm}(\mu, \nu) \mapsto C_p(\pi)$ is lower semi-continuous.
- $\text{Opt}_p(\mu, \nu)$ is a non-empty sequentially compact set.
- OT_p is lower semi-continuous, in the sense that for sequences $(\mu_n)_n, (\nu_n)_n$ in \mathcal{M}^p satisfying $\mu_n \xrightarrow{v} \mu$ and $\nu_n \xrightarrow{v} \nu$, we have

$$\text{OT}_p(\mu, \nu) \leq \liminf_{n \rightarrow \infty} \text{OT}_p(\mu_n, \nu_n).$$

Moreover, OT_p is a distance on \mathcal{M}^p .

These properties are mentioned in [FG10, pages 4-5] in the bounded case, and the corresponding proofs adapt straightforwardly to our framework (see Section 3.5).

REMARK 3.3. *If a (Borel) measure μ satisfies $\text{Pers}_p(\mu) < \infty$, then for any Borel set $A \subset \Omega$ satisfying $d(A, \partial\Omega) := \inf_{x \in A} d(x, \partial\Omega) > 0$, we have:*

$$\begin{aligned} \mu(A)d(A, \partial\Omega)^p &\leq \int_A d(x, \partial\Omega)^p d\mu(x) \\ &\leq \int_\Omega d(x, \partial\Omega)^p d\mu(x) = \text{Pers}_p(\mu) < \infty, \end{aligned} \tag{3.6}$$

so that $\mu(A) < \infty$. In particular, μ is automatically a Radon measure.

The following lemma gives a simple way to approximate a persistence measure (resp. diagram) with ones of finite masses.

LEMMA 3.4. *Let $\mu \in \mathcal{M}^p$. Fix $r > 0$, and let $A_r := \{x \in \Omega, d(x, \partial\Omega) \leq r\}$. Let $\mu^{(r)}$ be the restriction of μ to $\Omega \setminus A_r$. Then $\text{OT}_p(\mu^{(r)}, \mu) \rightarrow 0$ when $r \rightarrow 0$. Similarly, if $a \in \mathcal{D}^p$, we have $\mathbf{d}_p(a^{(r)}, a) \rightarrow 0$.*

Proof. Let $\pi \in \text{Adm}(\mu, \mu^{(r)})$ be the transport plan induced by the identity map on $\overline{\Omega} \setminus A_r$, and the projection onto $\partial\Omega$ on A_r . As π is sub-optimal, one has:

$$\text{OT}_p^p(\mu, \mu^{(r)}) \leq C_p(\pi) = \int_{A_r} d(x, \partial\Omega)^p d\mu(x) = \text{Pers}_p(\mu) - \text{Pers}_p(\mu^{(r)}).$$

Thus, by the monotone convergence theorem applied to μ with the functions $f_r : x \mapsto d(x, \partial\Omega)^p \cdot \mathbf{1}_{\Omega \setminus A_r}(x)$, we obtain that $\text{OT}_p(\mu, \mu^{(r)}) \rightarrow 0$ as $r \rightarrow 0$. Similar arguments show that $\mathbf{d}_p(a^{(r)}, a) \rightarrow 0$ as $r \rightarrow 0$. \square

The following result is central in this work: it shows that the metrics OT_p are extensions of the metrics \mathbf{d}_p .

PROPOSITION 3.5. *For $a, b \in \mathcal{D}^p$, $\text{OT}_p(a, b) = \mathbf{d}_p(a, b)$.*

Proof. Let $a, b \in \mathcal{D}^p$ be two persistence diagrams. Assume first that a, b have finitely many points $\{x_1 \dots x_n\}$ and $\{y_1 \dots y_m\}$ respectively. For the sake of simplicity, let us assume that each x_i and y_j has multiplicity one. Let γ be an optimal partial matching that achieves $\mathbf{d}_p(a, b)$ (see Section 2.1). It induces an admissible transport plan π by defining $\pi(x, y) = 1$ if $(x, y) \in \gamma$, and 0 otherwise. Furthermore,

$$\sum_{(x,y) \in \gamma} d(x, y)^p = \iint d(x, y)^p d\pi(x, y).$$

It yields $\mathbf{d}_p(a, b) \geq \text{OT}_p(a, b)$.

Conversely, an element $\pi \in \text{Adm}(a, b)$ that would achieve $\text{OT}_p(a, b)$ must be supported on $\{x_1 \dots x_n\} \cup \{\mathbf{p}_{\partial\Omega}(y_1) \dots \mathbf{p}_{\partial\Omega}(y_m)\} \times \{y_1 \dots y_m\} \cup \{\mathbf{p}_{\partial\Omega}(x_1) \dots \mathbf{p}_{\partial\Omega}(x_n)\}$, where $\mathbf{p}_{\partial\Omega}(x)$ denotes the orthogonal projection of $x \in \Omega$ onto the diagonal $\partial\Omega$. It can thus be described as a bi-stochastic matrix $P \in \mathbb{R}_+^{(n+m) \times (m+n)}$. If we denote by \mathcal{S}_{n+m} this set, it yields

$$\text{OT}_p(a, b) = \inf_{P \in \mathcal{S}_{n+m}} \langle C, P \rangle,$$

where $\langle C, P \rangle = \sum_{ij} C_{ij} P_{ij}$ (Frobenius dot product) and $C \in \mathbb{R}^{(n+m) \times (m+n)}$ is defined by

$$\begin{aligned} C_{ij} &= d(x_i, y_j) \text{ if } 1 \leq i \leq n, 1 \leq j \leq m, \\ C_{ij} &= d(x_i, \mathbf{p}_{\partial\Omega}(x_i)) \text{ if } 1 \leq i \leq n, j > m, \\ C_{ij} &= d(\mathbf{p}_{\partial\Omega}(y_j), y_j) \text{ if } i > n, 1 \leq j \leq m, \\ C_{ij} &= 0 \text{ if } i > n, j > m. \end{aligned}$$

By Birkhoff-Von Neumann theorem (Proposition 2.20), it follows that among the minimizers of $\langle C, P \rangle$ in \mathcal{S}_{n+m} , there must be a permutation matrix P^* . Such a permutation matrix induces a partial matching γ defined by $(x_i, y_j) \in \gamma \Leftrightarrow P_{ij}^* = 1$, while we have

$$\text{OT}_p(a, b) = \langle C, P^* \rangle = \sum_{(x,y) \in \gamma} d(x, y)^p \geq \mathbf{d}_p(a, b),$$

so that eventually $\text{OT}_p(a, b) = \mathbf{d}_p(a, b)$.

In the general case, let $r > 0$. Due to (3.6), the diagrams $a^{(r)}$ and $b^{(r)}$ defined in Lemma 3.4 have a finite mass (thus finite number of points). Therefore, $\mathbf{d}_p(a^{(r)}, b^{(r)}) = \text{OT}_p(a^{(r)}, b^{(r)})$. By Lemma 3.4, the former converges to $\mathbf{d}_p(a, b)$ while the latter converges to $\text{OT}_p(a, b)$, giving the conclusion. \square

As a consequence of this result, we will use OT_p to denote the distance between two elements of \mathcal{D}^p from now on.

PROPOSITION 3.6. *The space $(\mathcal{M}^p, \text{OT}_p)$ is a Polish space.*

As for Proposition 3.2, this proposition appears in [FG10, Proposition 2.7] in the bounded case, and its proof is straightforwardly adapted to our framework (see Section 3.5).

The following is one of the most important results of this chapter: a characterization of convergence in $(\mathcal{M}^p, \text{OT}_p)$.

THEOREM 3.7. *Let μ, μ_1, μ_2, \dots be measures in \mathcal{M}^p . Then,*

$$\text{OT}_p(\mu_n, \mu) \rightarrow 0 \Leftrightarrow \begin{cases} \mu_n \xrightarrow{v} \mu, \\ \text{Pers}_p(\mu_n) \rightarrow \text{Pers}_p(\mu). \end{cases} \quad (3.7)$$

This result is analog to the characterization of convergence of probability measures in the Wasserstein space (see [Vil08, Theorem 6.9]) and can be found in [FG10, Proposition 2.7] in the case where the ground space is bounded. While the proof of the direct implication can be easily adapted from [FG10], a new proof is needed for the converse implication. Both can be found in Section 3.5.

REMARK 3.8. *The assumption $\text{Pers}_p(\mu_n) \rightarrow \text{Pers}_p(\mu)$ is crucial to obtain OT_p -convergence assuming vague convergence. For instance, the sequence defined by $\mu_n := \delta_{(n, n+1)}$ converges vaguely to $\mu = 0$ and $(\text{Pers}_p(\mu_n))_n$ does converge (it is constant), while $\text{OT}_p(\mu_n, 0) \not\rightarrow 0$. This does not contradict Theorem 3.7 since $\text{Pers}_p(\mu) = 0 \neq \lim_n \text{Pers}_p(\mu_n)$.*

Theorem 3.7 implies some very useful results. First, it entails that the topology of the metric OT_p is stronger than the vague topology. As a consequence, the following corollary holds, using Proposition B.9 (\mathcal{D}^p is closed in \mathcal{M}^p for the vague topology).

COROLLARY 3.9. *\mathcal{D}^p is closed in \mathcal{M}^p for the metric OT_p .*

We recover in particular that the space of persistence diagrams $(\mathcal{D}^p, \text{OT}_p)$ is a Polish space (due to Proposition 3.6), a result already proved in [MMH11, Theorems 7 and 12] with a different approach.

Secondly, we show that the vague convergence of μ_n to μ along with the convergence of $\text{Pers}_p(\mu_n) \rightarrow \text{Pers}_p(\mu)$ is equivalent to the weak convergence of a weighted measure (see Appendix B for a definition of weak convergence, denoted by \xrightarrow{w} in the following). For $\mu \in \mathcal{M}^p$, let us introduce the Borel measure with finite mass $\mu^{[p]}$ defined, for a Borel subset $A \subset \Omega$, as:

$$\mu^{[p]}(A) = \int_A d(x, \partial\Omega)^p d\mu(x). \quad (3.8)$$

COROLLARY 3.10. *For a sequence $(\mu_n)_n$ and a persistence measure $\mu \in \mathcal{M}^p$, we have*

$$\text{OT}_p(\mu_n, \mu) \rightarrow 0 \text{ if and only if } \mu_n^{[p]} \xrightarrow{w} \mu^{[p]}.$$

Proof. Consider $\mu, \mu_1, \mu_2, \dots \in \mathcal{M}^p$ and assume that $\text{OT}_p(\mu_n, \mu) \rightarrow 0$. By Theorem 3.7, this is equivalent to $\mu_n \xrightarrow{v} \mu$ and $\mu_n^{[p]}(\Omega) = \text{Pers}_p(\mu_n) \rightarrow \text{Pers}_p(\mu) = \mu^{[p]}(\Omega)$. Since for any continuous function f compactly supported, the map $x \mapsto d(x, \partial\Omega)^p f(x)$ is also continuous and compactly supported, $\mu_n \xrightarrow{v} \mu$ implies $\mu_n^{[p]} \xrightarrow{v} \mu^{[p]}$. Likewise, the map $x \mapsto d(x, \partial\Omega)^{-p} f(x)$ is also continuous and compactly supported, so that $\mu_n^{(p)} \xrightarrow{v} \mu^{(p)}$ implies $\mu_n \xrightarrow{v} \mu$. Hence, $\mu_n \xrightarrow{v} \mu$ is equivalent to $\mu_n^{(p)} \xrightarrow{v} \mu^{(p)}$. By Proposition B.6, the vague convergence along with the convergence of the masses is also equivalent to the weak convergence $\mu_n^{[p]} \xrightarrow{w} \mu^{[p]}$. \square

We end this section with two results: a characterization of relatively compact sets in $(\mathcal{M}^p, \text{OT}_p)$, and a convexity property of OT_p^p that will be useful later on.

PROPOSITION 3.11. *A set F is relatively compact in $(\mathcal{M}^p, \text{OT}_p)$ if and only if the set $\{\mu^{[p]}, \mu \in F\}$ is tight and $\sup_{\mu \in F} \text{Pers}_p(\mu) < \infty$.*

Proof. From Corollary 3.10, the relative compactness of a set $F \subset \mathcal{M}^p$ for the metric OT_p is equivalent to the relative compactness of the set $\{\mu^{[p]}, \mu \in F\}$ for the weak convergence. Recall that all $\mu^{[p]}$ have a finite mass, as $\mu^{[p]}(\Omega) = \text{Pers}_p(\mu) < \infty$. Therefore, one can use Prokhorov's theorem (Proposition B.5) to conclude. \square

REMARK 3.12. *This characterization is equivalent to the one described in [MMH11, Theorem 21] for persistence diagrams. The notions introduced*

therein by the authors of off-diagonally birth-death boundedness, and uniformness are rephrased using the notion of tightness, standard in measure theory.

LEMMA 3.13. *For $1 \leq p < \infty$, the function $\text{OT}_p^p : \mathcal{M}^p \times \mathcal{M}^p \rightarrow \mathbb{R}$ is convex, that is for all $t \in [0, 1]$ and all $\mu_1, \mu_2, \nu_1, \nu_2 \in \mathcal{M}^p$, we have*

$$\text{OT}_p^p(t\mu_1 + (1-t)\mu_2, t\nu_1 + (1-t)\nu_2) \leq t\text{OT}_p^p(\mu_1, \nu_1) + (1-t)\text{OT}_p^p(\mu_2, \nu_2).$$

Proof. Fix $\mu_1, \mu_2, \nu_1, \nu_2 \in \mathcal{M}^p$ and $t \in [0, 1]$. Let $\pi_{11} \in \text{Opt}_p(\mu_1, \nu_1)$ and $\pi_{22} \in \text{Opt}_p(\mu_2, \nu_2)$. It is straightforward to check that $\pi := t\pi_{11} + (1-t)\pi_{22}$ is an admissible plan between $t\mu_1 + (1-t)\mu_2$ and $t\nu_1 + (1-t)\nu_2$. The cost of this admissible plan is $t\text{OT}_p^p(\mu_1, \nu_1) + (1-t)\text{OT}_p^p(\mu_2, \nu_2)$, which is therefore larger than $\text{OT}_p^p(t\mu_1 + (1-t)\mu_2, t\nu_1 + (1-t)\nu_2)$. \square

3.2 Persistence measures in the finite setting

In practice, many statistical results regarding persistence diagrams are stated for sets of diagrams with uniformly bounded number of points [KHN⁺15, CCO17], and the specific properties of OT_p in this setting are therefore of interest. Introduce for $m \geq 0$ the subset $\mathcal{M}_{\leq m}^p$ of \mathcal{M}^p defined as $\mathcal{M}_{\leq m}^p := \{\mu \in \mathcal{M}^p, \mu(\Omega) \leq m\}$, and the set \mathcal{M}_f^p of finite persistence measures, $\mathcal{M}_f^p := \bigcup_{m \geq 0} \mathcal{M}_{\leq m}^p$. Define similarly the set $\mathcal{D}_{\leq m}$ (resp. \mathcal{D}_f). Note that the assumption $\text{Pers}_p(a) < \infty$ is always satisfied for a finite diagram a (which is not true for general Radon measures), so that the exponent p is not needed when defining $\mathcal{D}_{\leq m}$ and \mathcal{D}_f .

PROPOSITION 3.14. *\mathcal{M}_f^p (resp. \mathcal{D}_f) is dense in \mathcal{M}^p (resp. \mathcal{D}^p) for the metric OT_p .*

Proof. This is a straightforward consequence of Lemma 3.4. Indeed, if $\mu \in \mathcal{M}^p$ and $r > 0$, then (3.6) implies that $\mu^{(r)}$ is of finite mass while $\mu^{(r)} \rightarrow \mu$ when $r \rightarrow 0$. \square

Let $\tilde{\Omega} = \Omega \sqcup \{\partial\Omega\}$ be the quotient of $\bar{\Omega}$ by the closed subset $\partial\Omega$ —i.e. we encode the diagonal by just one point (still denoted by $\partial\Omega$). The distance d on $\bar{\Omega}^2$ induces naturally a function \tilde{d} on $\tilde{\Omega}^2$, defined for $x, y \in \Omega$ by $\tilde{d}(x, y) = d(x, y)$, $\tilde{d}(x, \partial\Omega) = \tilde{d}(\partial\Omega, x) = d(x, \mathbf{p}_{\partial\Omega}(x))$ and $\tilde{d}(\partial\Omega, \partial\Omega) = 0$. However, \tilde{d} is not a distance since one can have $\tilde{d}(x, y) > \tilde{d}(x, \partial\Omega) + \tilde{d}(y, \partial\Omega)$. Define

$$\rho(x, y) := \min\{\tilde{d}(x, y), \tilde{d}(x, \partial\Omega) + \tilde{d}(y, \partial\Omega)\}. \quad (3.9)$$

It is straightforward to check that ρ is a distance on $\tilde{\Omega}$ and that $(\tilde{\Omega}, \rho)$ is a Polish space. One can then define the Wasserstein distance $W_{p,\rho}$ with respect to ρ for finite measures on $\tilde{\Omega}$ which have the same masses, that is the infimum of $\tilde{C}_p(\tilde{\pi}) := \iint_{\tilde{\Omega}^2} \rho(x, y)^p d\tilde{\pi}(x, y)$, for $\tilde{\pi}$ a transport plan with corresponding marginals (see Section 2.2.2). The following result states that the problem of computing the OT_p metric between two persistence measures with finite masses can be turned into the one of computing the Wasserstein distances between two measures supported on $\tilde{\Omega}$ with the same mass. For the sake of simplicity, we assume here that $d(x, y) = \|x - y\|_q$ with $q > 1$ so that for $x \in \Omega$, the set $\arg \min\{y \in \partial\Omega, d(x, y)\}$ is given by the orthogonal projection $\mathbf{p}_{\partial\Omega}(x)$ of x onto $\partial\Omega$. The following result could be seamlessly adapted to the case $q = 1$.

PROPOSITION 3.15. *Let $\mu, \nu \in \mathcal{M}_f^p$ and $r \geq \mu(\Omega) + \nu(\Omega)$. Define $\tilde{\mu} = \mu + (r - \mu(\Omega))\delta_{\partial\Omega}$ and $\tilde{\nu} = \nu + (r - \nu(\Omega))\delta_{\partial\Omega}$. Then $\text{OT}_p(\mu, \nu) = W_{p,\rho}(\tilde{\mu}, \tilde{\nu})$.*

Before proving Proposition 3.15, we need to introduce two lemmas. The first one essentially states that an element $\pi \in \text{Adm}(\mu, \nu)$ can be turned into an element $\tilde{\pi} \in \Pi(\tilde{\mu}, \tilde{\nu})$ such that π and $\tilde{\pi}$ have the same transport cost (for d and \tilde{d} respectively), and vice-versa. The second lemma gives a first inequality to prove Proposition 3.15. The proofs of these lemmas have been delayed to Section 3.5.

LEMMA 3.16. *Let $\mu, \nu \in \mathcal{M}_f^p$ and $r \geq \max(\mu(\Omega), \nu(\Omega))$. Let $\tilde{\mu} := \mu + (r - \mu(\Omega))\delta_{\partial\Omega}$, $\tilde{\nu} := \nu + (r - \nu(\Omega))\delta_{\partial\Omega}$ and $\mathbf{p}_{\partial\Omega} : \Omega \rightarrow \partial\Omega$ be the orthogonal projection on the diagonal.*

1. Define $T(\mu, \nu)$ the set of plans $\pi \in \text{Adm}(\mu, \nu)$ satisfying $\pi(\{(x, y) \in \Omega \times \partial\Omega, y \neq \mathbf{p}_{\partial\Omega}(x)\}) = \pi(\{(x, y) \in \partial\Omega \times \Omega, x \neq \mathbf{p}_{\partial\Omega}(y)\}) = 0$ along with $\pi(\partial\Omega \times \partial\Omega) = 0$. Then, $\text{Opt}_p(\mu, \nu) \subset T(\mu, \nu)$.
2. Let $\pi \in T(\mu, \nu)$ be such that $\mu(\Omega) + \pi(\partial\Omega \times \Omega) \leq r$. Define $\iota(\pi) \in \Pi(\tilde{\mu}, \tilde{\nu})$ by, for Borel sets $A, B \subset \Omega$,

$$\begin{cases} \iota(\pi)(A \times B) = \pi(A \times B), \\ \iota(\pi)(A \times \{\partial\Omega\}) = \pi(A \times \partial\Omega), \\ \iota(\pi)(\{\partial\Omega\} \times B) = \pi(\partial\Omega \times B), \\ \iota(\pi)(\{\partial\Omega\} \times \{\partial\Omega\}) = r - \mu(\Omega) - \pi(\partial\Omega \times \Omega) \geq 0. \end{cases} \quad (3.10)$$

Then, $C_p(\pi) = \iint_{\tilde{\Omega} \times \tilde{\Omega}} d(x, y)^p d\iota(\pi)(x, y)$.

3. Let $\tilde{\pi} \in \Pi(\tilde{\mu}, \tilde{\nu})$. Define $\kappa(\tilde{\pi}) \in T(\mu, \nu)$ by (where A, B are Borel sets),

$$\begin{cases} \kappa(\tilde{\pi})(A \times B) = \tilde{\pi}(A \times B), & A, B \subset \Omega, \\ \kappa(\tilde{\pi})(A \times B) = \tilde{\pi}((A \cap s^{-1}(B)) \times \{\partial\Omega\}), & A \subset \Omega, B \subset \partial\Omega, \\ \kappa(\tilde{\pi})(A \times B) = \tilde{\pi}(\{\partial\Omega\} \times (B \cap s^{-1}(A))), & A \subset \partial\Omega, B \subset \Omega, \\ \kappa(\tilde{\pi})(\partial\Omega, \partial\Omega) = 0. \end{cases}$$

Then, $\iint_{\tilde{\Omega} \times \tilde{\Omega}} d(x, y)^p d\tilde{\pi}(x, y) = C_p(\kappa(\tilde{\pi}))$.

LEMMA 3.17. Let $\mu, \nu \in \mathcal{M}_f^p$ and $r \geq \max(\mu(\Omega), \nu(\Omega))$. Let $\tilde{\mu} := \mu + (r - \mu(\Omega))\delta_{\partial\Omega}$, $\tilde{\nu} := \nu + (r - \nu(\Omega))\delta_{\partial\Omega}$. Then, $\text{OT}_p(\mu, \nu) \leq W_{p,\rho}(\tilde{\mu}, \tilde{\nu})$.

Proof of Proposition 3.15. Let $\pi \in T(\mu, \nu)$. As $\mu(\Omega) + \pi(\partial\Omega \times \Omega) \leq \mu(\Omega) + \nu(\Omega) \leq r$, one can define $\tilde{\pi} = \iota(\pi)$. Since $\rho(x, y) \leq d(x, y)$, we have $\tilde{C}_p(\tilde{\pi}) \leq \iint d(x, y)^p d\tilde{\pi}(x, y) = C_p(\pi)$ (Lemma 3.16). Taking the infimum gives $W_{p,\rho}(\tilde{\mu}, \tilde{\nu}) \leq \text{OT}_p(\mu, \nu)$. The other inequality holds according to Lemma 3.17. \square

REMARK 3.18. A careful read of the proof of Proposition 3.15 and relative lemmas shows that one could replace the metric ρ by

$$\rho_p : (x, y) \mapsto \min\{d(x, y), (d(x, \partial\Omega)^p + d(\partial\Omega, y)^p)^{\frac{1}{p}}\}.$$

This fact will be used in Section 3.4.

REMARK 3.19. The starting idea of this theorem—informally, “adding the mass of one diagram to the other and vice-versa”—is known in TDA as a bipartite graph matching [EH10, Ch. VIII.4] and used in practical computations [KMN17]. Here, Proposition 3.15 states that solving this bipartite graph matching problem can be formalized as computing a Wasserstein distance on the metric space $(\tilde{\Omega}, \rho)$ and, as such, makes sense (and remains true) for more general measures.

REMARK 3.20. Proposition 3.15 is useful for numerical purposes since it allows in applications, when dealing with a finite set of finite measures (in particular diagrams), to directly use the various tools developed in computational optimal transport [PC17] to compute Wasserstein distances. This alternative to the combinatorial algorithms considered in the literature [KMN17, TMMH14] will be presented in Chapter 5, based on [LCO18]. This result is also helpful to prove the existence of Fréchet means of sets of persistence measures (see Chapter 4).

3.3 The OT_∞ distance

In classical optimal transport, the ∞ -Wasserstein distance is known to have a much more erratic behavior than its $p < \infty$ counterparts [San15, Section 5.5.1]. However, in the context of persistence diagrams, the distance \mathbf{d}_∞ defined in Eq. (2.1) appears naturally as an interleaving distance between persistence modules (see Appendix A.2) and satisfies strong stability results: it is thus worthy of interest. It also happens that, when restricted to diagrams having some specific finiteness properties, most irregular behaviors are suppressed and a convenient characterization of convergence exists.

DEFINITION 3.21. *Let $\text{spt}(\mu)$ denote the support of a measure μ and*

$$\text{Pers}_\infty(\mu) := \sup\{d(x, \partial\Omega), x \in \text{spt}(\mu)\}.$$

Define

$$\mathcal{M}^\infty := \{\mu \in \mathcal{M}, \text{Pers}_\infty(\mu) < \infty\} \quad \text{and} \quad \mathcal{D}^\infty := \mathcal{D} \cap \mathcal{M}^\infty. \quad (3.11)$$

For $\mu, \nu \in \mathcal{M}^\infty$ and $\pi \in \text{Adm}(\mu, \nu)$, let

$$C_\infty(\pi) := \sup\{d(x, y), (x, y) \in \text{spt}(\pi)\}$$

and let

$$\text{OT}_\infty(\mu, \nu) := \inf_{\pi \in \text{Adm}(\mu, \nu)} C_\infty(\pi). \quad (3.12)$$

The set of transport plans minimizing (3.12) is denoted by $\text{Opt}_\infty(\mu, \nu)$.

PROPOSITION 3.22. *Let $\mu, \nu \in \mathcal{M}^\infty$. For the vague topology on E_Ω ,*

- *the map $\pi \in \text{Adm}(\mu, \nu) \mapsto C_\infty(\pi)$ is lower semi-continuous.*
- *The set $\text{Opt}_\infty(\mu, \nu)$ is a non-empty sequentially compact set.*
- *OT_∞ is lower semi-continuous.*

Moreover, OT_∞ is a distance on \mathcal{M}^∞ .

The proofs of these results are found in Section 3.5.

As in the case $p < \infty$, OT_∞ and \mathbf{d}_∞ coincide on \mathcal{D}^∞ .

PROPOSITION 3.23. *For $a, b \in \mathcal{D}^\infty$, $\text{OT}_\infty(a, b) = \mathbf{d}_\infty(a, b)$.*

Proof. Consider two diagrams $a, b \in \mathcal{D}^\infty$, written as $a = \sum_{i \in I} \delta_{x_i}$ and $b = \sum_{j \in J} \delta_{y_j}$, where $I, J \subset \mathbb{N}^*$ are (potentially infinite) sets of indices. The marginal constraints imply that a plan $\pi \in \text{Adm}(\mu, \nu)$ is supported on $(\{x_i\}_i \cup \partial\Omega) \times (\{y_j\}_j \cup \partial\Omega)$, and the cost of such a plan can always be reduced if some of the mass $\pi(\{x_i\}, \partial\Omega)$ (resp. $\pi(\partial\Omega, \{y_j\})$) is sent on a point other than the projection of x_i (resp. y_j) on the diagonal $\partial\Omega$. Introduce the matrix C indexed on $(-J \cup I) \times (-I \cup J)$ defined by

$$\begin{cases} C_{i,j} = d(x_i, y_j) & \text{for } i, j > 0, \\ C_{i,j} = d(\partial\Omega, y_j) & \text{for } i < 0, j > 0, \\ C_{i,j} = d(x_i, \partial\Omega) & \text{for } i > 0, j < 0, \\ C_{i,j} = 0 & \text{for } i, j < 0. \end{cases} \quad (3.13)$$

In this context, an element of $\text{Opt}(a, b)$ can be written a matrix P indexed on $(-J \cup I) \times (-I \cup J)$, and marginal constraints state that P must belong to the set of doubly stochastic matrices \mathcal{S} . Therefore, $\text{OT}_\infty(a, b) = \inf_{P \in \mathcal{S}} \sup\{C_{i,j}, (i, j) \in \text{spt}(P)\}$, where \mathcal{S} is the set of doubly stochastic matrices indexed on $(-J \cup I) \times (-I \cup J)$, and $\text{spt}(P)$ denotes the support of P , that is the set $\{(i, j), P_{i,j} > 0\}$.

Let $P \in \mathcal{S}$. For any $k \in \mathbb{N}$, and any set of distinct indices $\{i_1, \dots, i_k\} \subset -J \cup I$, we have

$$k = \sum_{k'=1}^k \underbrace{\sum_{j \in -I \cup J} P_{i_{k'}, j}}_{=1} = \sum_{j \in -I \cup J} \underbrace{\sum_{k'=1}^k P_{i_{k'}, j}}_{\leq 1}.$$

Thus, the cardinality of $\{j, \exists k' \text{ such that } (i_{k'}, j) \in \text{spt}(P)\}$ must be larger than k . Said differently, the marginal constraints impose that any set of k points in a must be matched to *at least* k points in b (points are counted with eventual repetitions here). Under such conditions, the Hall's marriage theorem (see [Hal86, p. 51]) guarantees the existence of a permutation matrix P' with $\text{spt}(P') \subset \text{spt}(P)$. As a consequence,

$$\begin{aligned} \sup\{C_{i,j}, (i, j) \in \text{spt}(P)\} &\geq \sup\{C_{i,j}, (i, j) \in \text{spt}(P')\} \\ &\geq \inf_{P' \in \mathcal{S}'} \sup\{C_{i,j}, (i, j) \in \text{spt}(P')\} = \mathbf{d}_\infty(a, b), \end{aligned}$$

where \mathcal{S}' denotes the set of permutations matrix indexed on $(-J \cup I) \times (-I \cup J)$. Taking the infimum on $P \in \mathcal{S}$ on the left-hand side and using that $\mathcal{S}' \subset \mathcal{S}$ finally gives that $\text{OT}_\infty(a, b) = \mathbf{d}_\infty(a, b)$. \square

PROPOSITION 3.24. *The space $(\mathcal{M}^\infty, \text{OT}_\infty)$ is complete.*

Proof. Let $(\mu_n)_n$ be a Cauchy sequence for OT_∞ . Fix a compact $K \subset \Omega$, and pick $\varepsilon = d(K, \partial\Omega)/2$. There exists n_0 such that for $n > n_0$, $\text{OT}_\infty(\mu_n, \mu_{n_0}) < \varepsilon$. Let $K_\varepsilon := \{x \in \Omega, d(x, K) \leq \varepsilon\}$. By considering $\pi_n \in \text{Opt}_\infty(\mu_n, \mu_{n_0})$, and since $\text{OT}_\infty(\mu_n, \mu_{n_0}) < \varepsilon$, we have that

$$\mu_n(K) = \pi_n(K \times \bar{\Omega}) = \pi_n(K \times K_\varepsilon) \leq \mu_{n_0}(K_\varepsilon). \quad (3.14)$$

Therefore, $(\mu_n(K))_n$ is uniformly bounded, and Proposition B.4 implies that $(\mu_n)_n$ is relatively compact. Finally, the exact same computations as in the proof of the completeness for $p < \infty$ (see Section 3.5) show that $(\mu_n)_n$ converges for the OT_∞ metric. \square

REMARK 3.25. *Unlike the case $p < \infty$, the space \mathcal{D}^∞ (and therefore \mathcal{M}^∞) is not separable. Indeed, for $I \subset \mathbb{N}$, define the diagram $a_I := \sum_{i \in I} \delta_{(i, i+1)} \in \mathcal{D}^\infty$. The family $\{a_I, I \subset \mathbb{N}\}$ is uncountable, and for two distinct I, I' , $\text{OT}_\infty(a_I, a_{I'}) = \frac{\sqrt{2}}{2}$. This result is similar to [BV18, Theorem 4.20].*

We now show that the direct implication in Theorem 3.7 still holds in the case $p = \infty$.

PROPOSITION 3.26. *Let μ, μ_1, μ_2, \dots be measures in \mathcal{M}^∞ . If $\text{OT}_\infty(\mu_n, \mu) \rightarrow 0$, then $(\mu_n)_n$ converges vaguely to μ .*

The proof is found in Section 3.5.

REMARK 3.27. *As for the case $1 \leq p < \infty$, Proposition 3.26 implies that OT_∞ metricizes the vague convergence, and thus using Proposition 3.23 and Proposition 3.28, we have that $(\mathcal{D}^\infty, d_\infty)$ is closed in $(\mathcal{M}^\infty, \text{OT}_\infty)$ and is—in particular—complete.*

In contrast to the case $p < \infty$, a converse implication for Proposition 3.26 does not hold, even on the subspace of persistence diagrams (see Figure 3.1). To recover a space with a structure more similar to the one of \mathcal{D}^p , it is useful to look at a smaller set. Introduce \mathcal{D}_0^∞ the set of persistence diagrams such that for all $r > 0$, there are a finite number of points of the diagram of persistence larger than r and recall that \mathcal{D}_f denotes the set of persistence diagrams with finite number of points.

PROPOSITION 3.28. *The closure of \mathcal{D}_f for the distance OT_∞ is \mathcal{D}_0^∞ .*

Proof. Consider $a \in \mathcal{D}_0^\infty$. By definition, for all $n \in \mathbb{N}$, a has a finite number of points with persistence larger than $\frac{1}{n}$, so that the restriction a_n of a to




Figure 3.1: Illustration of differences between OT_p , OT_∞ , and vague convergences. Blue color represents the mass on a point while red color designates distances. (a) A case where $\text{OT}_p(\mu_n, 0) \rightarrow 0$ for any $p < \infty$ while $\text{OT}_\infty(\mu_n, 0) = 1$. (b) A case where $\text{OT}_\infty(\mu_n, 0) \rightarrow 0$ while for all $p < \infty$, $\text{OT}_p(\mu_n, \mu) \rightarrow \infty$. (c) A sequence of persistence diagrams $a_n \in \mathcal{D}^\infty$, where $(a_n)_n$ converges vaguely to $a = \sum_i \delta_{x_i}$ and $\text{Pers}_\infty(a_n) = \text{Pers}_\infty(a)$, but (a_n) does not converge to a for OT_∞ .

points with persistence larger than $\frac{1}{n}$ belongs to \mathcal{D}_f . As $\text{OT}_\infty(a, a_n) \leq \frac{1}{n} \rightarrow 0$, \mathcal{D}_0^∞ is contained in the closure of \mathcal{D}_f .

Conversely, consider a diagram $a \in \mathcal{D}^\infty \setminus \mathcal{D}_0^\infty$. There is a constant $r > 0$ such that a has infinitely many points with persistence larger than r . For any finite diagram $a' \in \mathcal{D}_f$, we have $\text{OT}_\infty(a', a) \geq r$, thus a is not the limit for the OT_∞ metric of any sequence in \mathcal{D}_f . \square

REMARK 3.29. *The space \mathcal{D}_0^∞ is exactly the set introduced in [BGMP14, Theorem 3.5] as the completion of \mathcal{D}_f for the bottleneck metric d_∞ . Here, we recover that \mathcal{D}_0^∞ is complete as a closed subset of the complete space \mathcal{D}^∞ .*

Define for $r > 0$ and $a \in \mathcal{D}$, $a^{(r)}$ the persistence diagram restricted to $\{x \in \Omega, d(x, \partial\Omega) > r\}$ (as in Lemma 3.4). The following characterization of convergence holds in \mathcal{D}_0^∞ .

PROPOSITION 3.30. *Let a, a_1, a_2, \dots be persistence diagrams in \mathcal{D}_0^∞ . Then,*

$$\text{OT}_\infty(a_n, a) \rightarrow 0 \Leftrightarrow \begin{cases} a_n \xrightarrow{v} a, \\ (a_n^{(r)})_n \text{ is tight for all positive } r. \end{cases}$$

Proof of Proposition 3.30. Let us prove first the direct implication. Proposition 3.26 states that the convergence with respect to OT_∞ implies the vague convergence. Fix $r > 0$. By definition, $a^{(r)}$ is made of a finite number of points, all included in some open bounded set $U \subset \Omega$. As $a_n^{(r)}(U^c)$ is

a sequence of integers, the bottleneck convergence implies that for n large enough, $a_n^{(r)}(U^c)$ is equal to 0. Thus, $(a_n^{(r)})_n$ is tight.

Let us prove the converse. Consider $a \in \mathcal{D}_0^\infty$ and a sequence $(a_n)_n$ that converges vaguely to a , with $(a_n^{(r)})_n$ tight for all $r > 0$. Fix $r > 0$ and let x_1, \dots, x_K be an enumeration of the points in $a^{(r)}$, the point x_k being present with multiplicity $m_k \in \mathbb{N}$. Denote by $B(x, \varepsilon)$ (resp. $\bar{B}(x, \varepsilon)$) the open (resp. closed) ball of radius ε centered at x . By the Portmanteau theorem, for ε small enough,

$$\begin{cases} \liminf_{n \rightarrow \infty} a_n(B(x_k, \varepsilon)) \geq a(B(x_k, \varepsilon)) = m_k \\ \limsup_{n \rightarrow \infty} a_n(\bar{B}(x_k, \varepsilon)) \leq a(\bar{B}(x_k, \varepsilon)) = m_k, \end{cases}$$

so that, for n large enough, there are exactly m_k points of a_n in $B(x_k, \varepsilon)$ (since $(a_n(B(x_k, \varepsilon)))_n$ is a converging sequence of integers). The tightness of $(a_n^{(r)})_n$ implies the existence of some compact $K \subset \Omega$ such that for n large enough, $a_n^{(r)}(K^c) = 0$ (as the measures take their values in \mathbb{N}). Applying the Portmanteau theorem to the closed set $K' := K \setminus \bigcup_{i=1}^K B(x_i, \varepsilon)$ gives

$$\limsup_{n \rightarrow \infty} a_n^{(r)}(K') \leq a^{(r)}(K') = 0.$$

This implies that for n large enough, there are no other point in a_n with persistence larger than r and thus $\text{OT}_\infty(a^{(r)}, a_n)$ is less than or equal to $r + \varepsilon$. Finally,

$$\limsup_{n \rightarrow \infty} \text{OT}_\infty(a_n, a) \leq \limsup_{n \rightarrow \infty} \text{OT}_\infty(a_n, a^{(r)}) + r \leq 2r + \varepsilon.$$

Letting $\varepsilon \rightarrow 0$ then $r \rightarrow 0$, the bottleneck convergence holds. \square

As for the case of OT_p metrics, this allows us to characterize relatively compact subsets of \mathcal{D}_0^∞ . A similar result can be found in [PMK19, Theorem 11]. Here again,

COROLLARY 3.31. *A set $F \subset \mathcal{D}_0^\infty$ is relatively compact for the metric OT_∞ if and only if for all compact K , $\sup\{a(K), a \in F\}$ is finite and for all $r > 0$, there exists a compact K such that $a^r(\Omega \setminus K) = 0$ for all $a \in F$.*

3.4 Duality

We end this chapter by providing a duality result for the OT_p metrics we introduced. Our approach essentially mimics the standard proof done in standard optimal transport [Vil08, Ch. 5] and also builds on some elements introduced in [FG10].

To alleviate notations, let $c(x, y) := d(x, y)^p$ in the following, and denote by \mathcal{M}_c the set \mathcal{M}^p . Define also

$$\tilde{c}(x, y) := \min\{c(x, y), c(x, \partial\Omega) + c(\partial\Omega, y)\} = \rho_p^p,$$

where ρ_p was introduced in Remark 3.18. Introduce

$$\mathcal{A} := \{(x, y) \in \bar{\Omega}, \tilde{c}(x, y) = c(x, y)\}.$$

The following definitions are standard in convex optimization and are routinely used in optimal transport theory when it comes to study duality.

DEFINITION 3.32. For $\varphi : \bar{\Omega} \rightarrow \mathbb{R}$, define its c -transform as

$$\varphi^c(y) = \inf_x c(x, y) - \varphi(x).$$

We say that φ is c -concave if it is the c -transform of some function ψ , that is $\varphi = \psi^c$.

Define then the c -superdifferential of a c -concave function φ as the set

$$\partial^c \varphi := \{(x, y) \in \bar{\Omega} \times \bar{\Omega}, \varphi(x) + \varphi^c(y) = c(x, y)\}.$$

Finally, say that a subset X of $\bar{\Omega} \times \bar{\Omega}$ is c -cyclically monotone if for any $K \in \mathbb{N}$ and any finite sample $(x_1, y_1) \dots (x_K, y_K) \in X$, and any permutation $\sigma \in \mathfrak{S}_K$, one has

$$\sum_{i=1}^K c(x_i, y_i) \leq \sum_{i=1}^K c(x_i, y_{\sigma(i)}). \quad (3.15)$$

Notions of \tilde{c} -concavity, \tilde{c} -transform, \tilde{c} -superdifferential, \tilde{c} -cyclical monotonicity are defined similarly.

The following proposition is central to obtain a duality result.

PROPOSITION 3.33. Let $\mu, \nu \in \mathcal{M}_c$ and $\pi \in \text{Adm}(\mu, \nu)$. We have equivalence between

(i) $\pi \in \text{Opt}(\mu, \nu)$.

(ii) $\text{spt}(\pi) \subset \mathcal{A}$ and $\text{spt}(\pi)$ is \tilde{c} -cyclically monotone.

(iii) There exists a c -concave function φ such that φ and φ^c both vanish on $\partial\Omega$ and $\text{spt}(\pi) \subset \partial^c \varphi$.

As a corollary, if $\mu_n \xrightarrow{v} \mu$ and $\nu_n \xrightarrow{v} \nu$, and $\pi_n \in \text{Opt}(\mu_n, \nu_n)$, then $\pi_n \xrightarrow{v} \pi \in \text{Opt}(\mu, \nu)$.

The proof essentially follows [FG10] and can be found in Section 3.5.

THEOREM 3.34 (Duality). *One has*

$$L_c(\mu, \nu) := \inf_{\pi \in \text{Adm}(\mu, \nu)} \langle c, \pi \rangle = \sup_{f, g \in \Psi_c} \langle f, \mu \rangle + \langle g, \nu \rangle \quad (3.16)$$

where Ψ_c denotes the set of continuous functions $f, g : \bar{\Omega} \rightarrow \mathbb{R}$ satisfying

- $\forall x, y \in \bar{\Omega}, f(x) + g(y) \leq c(x, y)$.
- $f(\partial\Omega) = g(\partial\Omega) = 0$

Proof. The proof relies on the equivalence between (i) and (iii) in Proposition 3.33. Let φ be a continuous function supported on $\bar{\Omega}$ that vanishes on the diagonal $\partial\Omega$. By definition of the c -transform, for all $x, y \in \bar{\Omega}$, one has $\varphi(x) + \varphi^c(y) \leq c(x, y)$. Therefore, for any $\pi \in \text{Adm}(\mu, \nu)$, we have

$$\int_{\Omega} \varphi(x) d\mu(x) + \int_{\Omega} \varphi^c(y) d\nu(y) \leq \iint_{\bar{\Omega} \times \bar{\Omega}} c(x, y) d\pi(x, y)$$

thanks to the marginal constraints (and that φ, φ^c both vanish on the diagonal). Thus,

$$\sup_{\varphi} \langle \varphi, \mu \rangle + \langle \varphi^c, \nu \rangle \leq \inf_{\pi \in \text{Adm}(\mu, \nu)} \langle c, \pi \rangle, \quad (3.17)$$

where the supremum in the left hand side is taken over the continuous functions φ that vanish on the diagonal. This inequality is known as a *duality gap*.

Now, consider $\pi \in \text{Opt}(\mu, \nu)$. We know thanks to (iii) that one can pick φ such that $\text{spt}(\pi) \subset \partial^c \varphi$, which gives that $\langle \varphi, \mu \rangle + \langle \varphi^c, \nu \rangle = \langle c, \pi \rangle = \text{OT}_p(\mu, \nu)$. Therefore, it follows that the duality gap in (3.17) is actually 0.

Finally, to prove the claimed result, it suffices to observe that by definition of the c -transform, one has $g(y) \leq \varphi^c(y)$ for any function g vanishing on the diagonal that satisfies $\varphi(x) + g(y) \leq c(x, y)$, and thus (changing φ for f), we obtain the inequality. \square

These results have some nice consequences in the particular case $p = 1$, that is $c(x, y) = d(x, y)$. One can observe (see [San15, Prop. 3.1]) that in this case, a function $\varphi : \bar{\Omega} \rightarrow \mathbb{R}$ is c -concave if and only if it is Lipschitz continuous with constant less than 1 (with respect to d), and for such a function φ , one has $\varphi^c = -\varphi$. Thus, we obtain the following Kantorovich-Rubinstein duality result:

PROPOSITION 3.35. *Let $p = 1$, that is $c(x, y) = d(x, y) = \|x - y\|$, and $\mu, \nu \in \mathcal{M}^1$. We have*

$$\text{OT}_1(\mu, \nu) = \max_{f \in \mathcal{L}} (\mu - \nu)(f), \quad (3.18)$$

where \mathcal{L} denotes the set of Lipschitz continuous functions $f : \overline{\Omega} \rightarrow \mathbb{R}$ with Lipschitz constant less than 1 and that satisfy $f(\partial\Omega) = 0$.

Note that a similar duality result for finite persistence diagram metrics was proved in [BE19].

3.5 Proofs

Proofs of Proposition 3.2 and Proposition 3.22.

- For $\pi \in \text{Adm}(\mu, \nu)$ supported on E_Ω , and for any compact sets $K, K' \subset \Omega$, one has $\pi((K \times \overline{\Omega}) \cup (\overline{\Omega} \times K')) \leq \mu(K) + \nu(K') < \infty$. As any compact subset of E_Ω is included in a set of the form $(K \times \overline{\Omega}) \cup (\overline{\Omega} \times K')$, Proposition B.4 implies that $\text{Adm}(\mu, \nu)$ is relatively compact for the vague convergence on E_Ω . Also, if a sequence $(\pi_n)_n$ in $\text{Adm}(\mu, \nu)$ converges vaguely to some $\pi \in \mathcal{M}(E_\Omega)$, then the marginals of π on Ω are still μ and ν . Indeed, if f is a continuous function with compact support on Ω , then

$$\begin{aligned} \int_{E_\Omega} f(x) d\pi(x, y) &= \lim_n \int_{E_\Omega} f(x) d\pi_n(x, y) \\ &= \lim_n \int_{\Omega} f(x) d\mu_n(x) \\ &= \int_{\Omega} f(x) d\mu(x), \end{aligned}$$

and we show likewise that the second marginal of π is ν . Thus, $\text{Adm}(\mu, \nu)$ is closed in $\mathcal{M}(E_\Omega)$: it is therefore sequentially compact.

- To prove the second point of Proposition 3.2, consider π, π_1, π_2, \dots such that $\pi_n \xrightarrow{v} \pi$, and introduce $\pi'_n : A \mapsto \iint_A d(x, y)^p d\pi_n$. The sequence $(\pi'_n)_n$ converges vaguely to $\pi' : A \mapsto \iint_A d(x, y)^p d\pi$. The Portmanteau theorem (Proposition B.7) applied with the open set E_Ω to the measures $\pi'_n \xrightarrow{v} \pi'$ implies that $C_p(\pi) = \pi'(E_\Omega) \leq \liminf_n \pi'_n(E_\Omega) = \liminf_n C_p(\pi_n)$, i.e. C_p is lower semi-continuous.

- We now prove the lower semi-continuity of C_∞ . Let $(\pi_n)_n$ be a sequence converging vaguely to π on E_Ω and let $r > \liminf_{n \rightarrow \infty} C_\infty(\pi_n)$. The set $U_r = \{(x, y) \in E_\Omega, d(x, y) > r\}$ is open. By the Portmanteau theorem (Proposition B.7), we have $0 = \liminf_{n \rightarrow \infty} \pi_n(U_r) \geq \pi(U_r)$. Therefore, $\text{spt}(\pi) \subset U_r^c$ and $C_\infty(\pi) \leq r$. As this holds for any $r > \liminf_{n \rightarrow \infty} C_\infty(\pi_n)$, we have $\liminf_{n \rightarrow \infty} C_\infty(\pi_n) \geq C_\infty(\pi)$.
- We show that for any $1 \leq p \leq \infty$, the lower semi-continuity of C_p and the sequential compactness of $\text{Adm}(\mu, \nu)$ imply that (1.) $\text{Opt}_p(\mu, \nu)$ is a non-empty compact set for the vague topology on E_Ω and that (2.) OT_p is lower semi-continuous.
 1. Let $(\pi_n)_n$ be a minimizing sequence of (2.22) in $\text{Adm}(\mu, \nu)$. As $\text{Adm}(\mu, \nu)$ is sequentially compact, it has an adherence value π , and the lower semi-continuity implies that

$$C_p(\pi) \leq \liminf_{n \rightarrow \infty} C_p(\pi_n) = \text{OT}_p^p(\mu, \nu),$$

so that $\text{Opt}_p(\mu, \nu)$ is non-empty. Using once again the lower semi-continuity of C_p , if a sequence in $\text{Opt}_p(\mu, \nu)$ converges to some limit, then the cost of the limit is smaller than (and thus equal to) $\text{OT}_p^p(\mu, \nu)$, i.e. the limit is in $\text{Opt}_p(\mu, \nu)$. The set $\text{Opt}_p(\mu, \nu)$ being closed in the sequentially compact set $\text{Adm}(\mu, \nu)$, it is also sequentially compact.

2. Let $\mu_n \xrightarrow{v} \mu$ and $\nu_n \xrightarrow{v} \nu$. One has $\liminf_n \text{OT}_p(\mu_n, \nu_n) = \lim_k \text{OT}_p(\mu_{n_k}, \nu_{n_k})$ for some subsequence $(n_k)_k$. For ease of notation, we will still use the index n to denote this subsequence. If the limit is infinite, there is nothing to prove. Otherwise, consider $\pi_n \in \text{Opt}_p(\mu_n, \nu_n)$. For any compact sets $K, K' \subset \Omega$, one has $\pi_n((K \times \bar{\Omega}) \cup (\bar{\Omega} \times K')) \leq \sup_n \mu_n(K) + \sup_n \nu_n(K') < \infty$. Therefore, by Proposition B.4, there exists a subsequence $(\pi_{n_k})_k$ which converges vaguely to some measure $\pi \in \text{Adm}(\mu, \nu)$. Note that the first (resp. second) marginal of π is equal to the limit μ (resp. ν) of the first (resp. second) marginal of $(\pi_{n_k})_k$, so that π is in $\text{Adm}(\mu, \nu)$. Therefore,

$$\text{OT}_p^p(\mu, \nu) \leq C_p(\pi) \leq \liminf_{n \rightarrow \infty} C_p(\pi_n) = \liminf_{n \rightarrow \infty} \text{OT}_p^p(\mu_n, \nu_n).$$

- Finally, we prove that OT_p is a metric on \mathcal{M}^p . Let $\mu, \nu, \lambda \in \mathcal{M}^p$. The symmetry of OT_p is clear. If $\text{OT}_p(\mu, \nu) = 0$, then there exists $\pi \in \text{Adm}(\mu, \nu)$ supported on $\{(x, x), x \in \Omega\}$. Therefore, for a Borel

set $A \subset \Omega$, $\mu(A) = \pi(A \times \overline{\Omega}) = \pi(A \times A) = \pi(\overline{\Omega} \times A) = \nu(A)$, and $\mu = \nu$. To prove the triangle inequality, we need a variant on the gluing lemma, stated in [FG10, Lemma 2.1]: for $\pi_{12} \in \text{Opt}(\mu, \nu)$ and $\pi_{23} \in \text{Opt}(\nu, \lambda)$ there exists a measure $\gamma \in \mathcal{M}(\overline{\Omega}^3)$ such that the marginal corresponding to the first two entries (resp. two last entries), when restricted to E_Ω , is equal to π_{12} (resp. π_{23}) and is concentrated on $\{(x, x), x \in \partial\Omega\}$ on $\partial\Omega \times \partial\Omega$. Note that the corresponding transportation cost on $\partial\Omega \times \partial\Omega$ is 0. Note also that the measure on $\overline{\Omega} \times \overline{\Omega}$ defined by $\int_y d\gamma(x, y, z)$ has μ and λ as marginals (when restricted to Ω), that is it belongs to $\text{Adm}(\mu, \lambda)$. Therefore, by the triangle and Minkowski inequality,

$$\begin{aligned} \text{OT}_p(\mu, \lambda) &\leq \left(\int_{\overline{\Omega}^3} d(x, z)^p d\gamma(x, y, z) \right)^{1/p} \\ &\leq \left(\int_{\overline{\Omega}^3} d(x, y)^p d\gamma(x, y, z) \right)^{1/p} + \left(\int_{\overline{\Omega}^3} d(y, z)^p d\gamma(x, y, z) \right)^{1/p} \\ &= \left(\int_{\overline{\Omega}^2} d(x, y)^p d\pi_{12}(x, y) \right)^{1/p} + \left(\int_{\overline{\Omega}^2} d(y, z)^p d\pi_{23}(y, z) \right)^{1/p} \\ &= \text{OT}_p(\mu, \nu) + \text{OT}_p(\nu, \lambda). \end{aligned}$$

The proof is similar for $p = \infty$.

□

Proof of Proposition 3.6. We first show the separability. Consider for $k > 0$ a partition of Ω into squares (C_i^k) of side length 2^{-k} , centered at points x_i^k . Let F be the set of all measures of the form $\sum_{i \in I} q_i \delta_{x_i^k}$ for q_i positive rationals, $k > 0$ and I a finite subset of \mathbb{N} . Our goal is to show that the countable set F is dense in \mathcal{M}^p . Fix $\varepsilon > 0$, and $\mu \in \mathcal{M}^p$. The proof is in three steps.

1. Since $\text{Pers}_p(\mu) < \infty$, there exists a compact $K \subset \Omega$ such that $\text{Pers}_p(\mu) - \text{Pers}_p(\mu_0) < \varepsilon^p$, where μ_0 is the restriction of μ to K . By considering the transport plan between μ and μ_0 induced by the identity map on K and the projection onto the diagonal on $\overline{\Omega} \setminus K$, it follows that $\text{OT}_p^p(\mu, \mu_0) \leq \text{Pers}_p(\mu) - \text{Pers}_p(\mu_0) \leq \varepsilon^p$.
2. Consider k such that $2^{-k} \leq \varepsilon / (\sqrt{2} \mu(K)^{1/p})$ and denote by I the indices corresponding to squares C_i^k intersecting K . Let us introduce $\mu_1 = \sum_{i \in I} \mu_0(C_i^k) \delta_{x_i^k}$. One can create a transport map between μ_0

and μ_1 by mapping each square C_i^k to its center x_i^k , so that

$$\text{OT}_p(\mu_0, \mu_1) \leq \left(\sum_i \mu_0(C_i^k) (\sqrt{2} \cdot 2^{-k})^p \right)^{1/p} \leq \mu(K)^{1/p} \sqrt{2} \cdot 2^{-k} \leq \varepsilon.$$

3. Consider, for $i \in I$, q_i a rational number satisfying $q_i \leq \mu_0(C_i^k)$ and $|\mu_0(C_i^k) - q_i| \leq \varepsilon^p / (\sum_{i \in I} d(x_i^k, \partial\Omega)^p)$. Let $\mu_2 = \sum_{i \in I} q_i \delta_{x_i^k}$. Consider the transport plan between μ_2 and μ_1 that fully transports μ_2 onto μ_1 , and transport the remaining mass in μ_1 onto the diagonal. Then,

$$\text{OT}_p(\mu_1, \mu_2) \leq \left(\sum_{i \in I} |\mu_0(C_i^k) - q_i| d(x_i^k, \partial\Omega)^p \right)^{1/p} \leq \varepsilon.$$

As $\mu_2 \in F$ and $\text{OT}_p(\mu, \mu_2) \leq 3\varepsilon$, the separability is proven.

To prove that the space is complete, consider a Cauchy sequence $(\mu_n)_n$. As $(\text{Pers}_p(\mu_n))_n = (\text{OT}_p^p(\mu_n, 0))_n$ is a Cauchy sequence, it is bounded. Therefore, for $K \subset \Omega$ a compact set, (3.6) implies that $\sup_n \mu_n(K) < \infty$. Proposition B.4 implies that $(\mu_n)_n$ is relatively compact for the vague topology on Ω . Consider $(\mu_{n_k})_k$ a subsequence converging vaguely on Ω to some measure μ . By the lower semi-continuity of OT_p ,

$$\text{Pers}_p(\mu) = \text{OT}_p^p(\mu, 0) \leq \liminf_{k \rightarrow \infty} \text{OT}_p^p(\mu_{n_k}, 0) < \infty,$$

so that $\mu \in \mathcal{M}^p$. Using once again the lower semi-continuity of OT_p ,

$$\begin{aligned} \text{OT}_p(\mu_n, \mu) &\leq \liminf_{k \rightarrow \infty} \text{OT}_p(\mu_n, \mu_{n_k}) \\ \lim_{n \rightarrow \infty} \text{OT}_p(\mu_n, \mu) &\leq \lim_{n \rightarrow \infty} \liminf_{k \rightarrow \infty} \text{OT}_p(\mu_n, \mu_{n_k}) = 0, \end{aligned}$$

ensuring that $\text{OT}_p(\mu_n, \mu) \rightarrow 0$, that is the space is complete. \square

Proof of the direct implication of Theorem 3.7. Let μ, μ_1, μ_2, \dots be elements of \mathcal{M}^p and assume that the sequence $(\text{OT}_p(\mu_n, \mu))_n$ converges to 0. The triangle inequality implies that $\text{Pers}_p(\mu_n) = \text{OT}_p^p(\mu_n, 0)$ converges to $\text{Pers}_p(\mu) = \text{OT}_p^p(\mu, 0)$. Let $f \in C_c(\Omega)$, whose support is included in some compact set K . For any $\varepsilon > 0$, there exists a Lipschitz function f_ε , with Lipschitz constant L and whose support is included in K , with the ∞ -norm $\|f - f_\varepsilon\|_\infty$ less than or equal to ε . The convergence of $\text{Pers}_p(\mu_n)$ and (3.6) imply that $\sup_k \mu_k(K) < \infty$. Let $\pi_n \in \text{Opt}_p(\mu_n, \mu)$, we have

$$\begin{aligned} |\mu_n(f) - \mu(f)| &\leq |\mu_n(f - f_\varepsilon)| + |\mu(f - f_\varepsilon)| + |\mu_n(f_\varepsilon) - \mu(f_\varepsilon)| \\ &\leq (\mu_n(K) + \mu(K))\varepsilon + |\mu_n(f_\varepsilon) - \mu(f_\varepsilon)| \\ &\leq (\sup_k \mu_k(K) + \mu(K))\varepsilon + |\mu_n(f_\varepsilon) - \mu(f_\varepsilon)|. \end{aligned}$$

Also,

$$\begin{aligned}
|\mu_n(f_\varepsilon) - \mu(f_\varepsilon)| &\leq \iint_{\overline{\Omega}^2} |f_\varepsilon(x) - f_\varepsilon(y)| d\pi_n(x, y) \quad \text{where } \pi_n \in \text{Opt}(\mu_n, \mu) \\
&\leq L \iint_{(K \times \overline{\Omega}) \cup (\overline{\Omega} \times K)} d(x, y) d\pi_n(x, y) \\
&\leq L \pi_n((K \times \overline{\Omega}) \cup (\overline{\Omega} \times K))^{1-\frac{1}{p}} \left(\iint_{(K \times \overline{\Omega}) \cup (\overline{\Omega} \times K)} d(x, y)^p d\pi_n(x, y) \right)^{\frac{1}{p}} \\
&\quad \text{by Hölder's inequality.} \\
&\leq L \left(\sup_k \mu_k(K) + \mu(K) \right)^{1-\frac{1}{p}} \text{OT}_p(\mu_n, \mu) \rightarrow 0.
\end{aligned}$$

This last quantity converges to 0 as n goes to ∞ for fixed ε . Therefore, taking the limsup in n and then letting $\varepsilon \rightarrow 0$, we obtain that $\mu_n(f) \rightarrow \mu(f)$. \square

Proof of the converse implication of Theorem 3.7. Let $\mu, \mu_1, \mu_2 \dots$ be elements of \mathcal{M}^p and assume that $\mu_n \xrightarrow{v} \mu$ and $\text{Pers}_p(\mu_n) \rightarrow \text{Pers}_p(\mu)$. Since

$$\text{OT}_p(\mu_n, \mu) \leq \text{OT}_p(\mu_n, 0) + \text{OT}_p(\mu, 0) = \text{Pers}_p(\mu_n)^{1/p} + \text{Pers}_p(\mu)^{1/p},$$

the sequence $(\text{OT}_p(\mu_n, \mu))_n$ is bounded. Thus, if we show that $(\text{OT}_p(\mu_n, \mu))_n$ admits 0 as unique accumulation point, then the convergence holds. Up to extracting a subsequence, we may assume that $(\text{OT}_p(\mu_n, \mu))_n$ converges to some limit. Let $(\pi_n)_n \in \text{Opt}(\mu_n, \mu)^\mathbb{N}$ be the corresponding optimal transport plans. Let K be a compact subset of Ω . Recall (Proposition B.4 in Appendix B) that relative compactness for the vague convergence of a sequence $(\mu_n)_n$ is equivalent to $\sup_n \{\mu_n(K)\} < \infty$ for every compact $K \subset \Omega$. Therefore, for any compact $K \subset \Omega$, and $n \in \mathbb{N}$,

$$\pi_n((K \times \Omega) \cup (\Omega \times K)) \leq \mu_n(K) + \mu(K) \leq \sup_k \mu_k(K) + \mu(K) < \infty.$$

As any compact of E_Ω is included in some set of the form $(K \cup \Omega) \times (\Omega \cup K)$, for $K \subset \Omega$ compact, using Proposition B.4 again, it follows that $(\pi_n)_n$ is also relatively compact for the vague topology.

Let thus π be the limit of any converging subsequence of $(\pi_n)_n$, which indexes are still denoted by n . As $\mu_n \xrightarrow{v} \mu$, π is necessarily in $\text{Opt}_p(\mu, \mu)$ (see [FG10, Prop. 2.3]), i.e. π is supported on $\{(x, x), x \in \Omega\}$. The vague convergence of $(\mu_n)_n$ and the convergence of $(\text{Pers}_p(\mu_n))_n$ to $\text{Pers}_p(\mu)$ imply

that for a given compact set $K \subset \Omega$, whose complementary set in Ω is denoted by K^c , its interior set is denoted by $\overset{\circ}{K}$, and its boundary by ∂K we have

$$\begin{aligned}
& \limsup_{n \rightarrow \infty} \int_{K^c} d(x, \partial\Omega)^p d\mu_n(x) \\
&= \limsup_{n \rightarrow \infty} \left(\text{Pers}_p(\mu_n) - \int_K d(x, \partial\Omega)^p d\mu_n(x) \right) \\
&= \lim_n \text{Pers}_p(\mu) - \liminf_n \int_{\overset{\circ}{K}} d(x, \partial\Omega)^p d\mu_n(x) - \liminf_n \int_{\partial K} d(x, \partial\Omega)^p d\mu_n(x) \\
&\leq \text{Pers}_p(\mu) - \int_{\overset{\circ}{K}} d(x, \partial\Omega)^p d\mu(x) \text{ by the Portmanteau theorem} \\
&= \int_{K^c} d(x, \partial\Omega)^p d\mu(x).
\end{aligned}$$

Recall that $\text{Pers}_p(\mu) < \infty$. Therefore, for $\varepsilon > 0$, there exists some compact set $K \subset \Omega$, with

$$\limsup_n \int_{K^c} d(x, \partial\Omega)^p d\mu_n(x) < \varepsilon \quad \text{and} \quad \int_{K^c} d(x, \partial\Omega)^p d\mu(x) < \varepsilon. \quad (3.19)$$

Let $\mathbf{p}_{\partial\Omega} : \Omega \rightarrow \partial\Omega$ be the orthogonal projection on $\partial\Omega$. We consider the following transport plan $\tilde{\pi}_n$ (consider informally that what went from K to K^c and from K^c to K is now transported onto the diagonal, while everything else is unchanged):

$$\begin{cases}
\tilde{\pi}_n = \pi_n & \text{on } K^2 \sqcup (K^c)^2, \\
\tilde{\pi}_n = 0 & \text{on } K \times K^c \sqcup K^c \times K, \\
\tilde{\pi}_n(A \times B) = \pi_n(A \times B) + \pi_n(A \times (\mathbf{p}_{\partial\Omega}^{-1}(B) \cap K^c)) & \text{for } A \subset K, B \subset \partial\Omega, \\
\tilde{\pi}_n(A \times B) = \pi_n(A \times B) + \pi_n(A \times (\mathbf{p}_{\partial\Omega}^{-1}(B) \cap K)) & \text{for } A \subset K^c, B \subset \partial\Omega, \\
\tilde{\pi}_n(A \times B) = \pi_n(A \times B) + \pi_n((\mathbf{p}_{\partial\Omega}^{-1}(A) \cap K^c) \times B) & \text{for } A \subset \partial\Omega, B \subset K, \\
\tilde{\pi}_n(A \times B) = \pi_n(A \times B) + \pi_n((\mathbf{p}_{\partial\Omega}^{-1}(A) \cap K) \times B) & \text{for } A \subset \partial\Omega, B \subset K^c.
\end{cases} \quad (3.20)$$

Note that $\tilde{\pi}_n \in \text{Adm}(\mu_n, \mu)$: for instance, for $A \subset K$ a Borel set,

$$\begin{aligned}
\tilde{\pi}_n(A \times \overline{\Omega}) &= \tilde{\pi}_n(A \times K) + \tilde{\pi}_n(A \times K^c) + \tilde{\pi}_n(A \times \partial\Omega) \\
&= \pi_n(A \times K) + 0 + \pi_n(A \times \partial\Omega) + \pi_n(A \times (\mathbf{p}_{\partial\Omega}^{-1}(\partial\Omega) \cap K^c)) \\
&= \pi_n(A \times \overline{\Omega}) = \mu_n(A),
\end{aligned}$$

and it is shown likewise that the other constraints are satisfied. As $\tilde{\pi}_n$ is suboptimal, $\text{OT}_p^p(\mu_n, \mu) \leq \int_{\overline{\Omega}^2} d(x, y)^p d\tilde{\pi}_n(x, y)$. The latter integral is equal to a sum of different terms, and we will show that each of them converges to 0. Assume without loss of generality that the compact set K belongs to an increasing sequence of compact sets whose union is Ω , with $\pi(\partial(K \times K)) = 0$ for all compacts of the sequence.

- We have $\iint_{K^2} d(x, y)^p d\tilde{\pi}_n(x, y) = \iint_{K^2} d(x, y)^p d\pi_n(x, y)$. The lim sup of the integral is less than or equal to $\iint_{K^2} d(x, y)^p d\pi(x, y)$ by the Portmanteau theorem (applied to the sequence $(d(x, y)^p d\pi_n(x, y))_n$), and, recalling that π is supported on the diagonal of E_Ω , this integral is equal to 0.
- For optimality reasons, any optimal transport plan must be supported on $\{d(x, y)^p \leq d(x, \partial\Omega)^p + d(y, \partial\Omega)^p\}$ (this fact is detailed in [FG10, Prop. 2.3]). It follows that

$$\begin{aligned} \iint_{(K^c)^2} d(x, y)^p d\tilde{\pi}_n(x, y) &= \iint_{(K^c)^2} d(x, y)^p d\pi_n(x, y) \\ &\leq \int_{K^c} d(x, \partial\Omega)^p d\mu_n(x) + \int_{K^c} d(y, \partial\Omega)^p d\mu(y). \end{aligned}$$

Taking the lim sup in n , and then letting K go to Ω , this quantity converges to 0 by (3.19).

- We have

$$\iint_{K \times \partial\Omega} d(x, \partial\Omega)^p d\tilde{\pi}_n(x, y) = \int_K d(x, \partial\Omega)^p d\mu_n(x) - \iint_{K^2} d(x, \partial\Omega)^p d\pi_n(x, y).$$

By the Portmanteau theorem applied to the sequence $(d(x, \partial\Omega)^p d\mu_n(x))_n$, we obtain that

$$\limsup_n \int_K d(x, \partial\Omega)^p d\mu_n(x) \leq \int_K d(x, \partial\Omega)^p d\mu(x).$$

Recall that we assume that $\pi(\partial(K \times K)) = 0$. By applying the second characterization of the Portmanteau theorem on the second term to the sequence $(d(x, y)^p d\pi_n(x, y))_n$, and using that π is supported on the diagonal of E_Ω , we have

$$\begin{aligned} \limsup_n - \iint_{K^2} d(x, \partial\Omega)^p d\pi_n(x, y) &\leq - \iint_{K^2} d(x, \partial\Omega)^p d\pi(x, y) \\ &= - \int_K d(x, \partial\Omega)^p d\mu(x) \end{aligned}$$

as $\pi(\partial(K \times K)) = 0$. Therefore, the lim sup of this (non-negative) integral is equal to 0.

- The three remaining terms (corresponding to the three last lines of the definition (3.20)) are treated likewise.

Finally, we have proven that $(\text{OT}_p(\mu_n, \mu))_n$ is bounded and that for any converging subsequence $(\mu_{n_k})_k$, $\text{OT}_p(\mu_{n_k}, \mu)$ converges to 0. It follows that $\text{OT}_p(\mu_n, \mu) \rightarrow 0$. \square

Proof of Lemma 3.16.

1. Consider $\pi \in \text{Adm}(\mu, \nu)$, and define π' that coincides with π on $\Omega \times \Omega$, and is such that we enforce that the mass located at some $x \in \Omega$ transported on the diagonal to be transported on its orthogonal projection $\mathbf{p}_{\partial\Omega}(x)$: more precisely, for all Borel set $A \subset \Omega$, $B \subset \partial\Omega$, $\pi'(A \times B) = \pi((\mathbf{p}_{\partial\Omega}^{-1}(B) \cap A) \times B)$ and $\pi'(B \times A) = \pi(B \times (\mathbf{p}_{\partial\Omega}^{-1}(B) \cap A))$. Note that $\pi' \in T(\mu, \nu)$. Since $\mathbf{p}_{\partial\Omega}(x)$ is the unique minimizer of $y \mapsto d(x, y)^p$, it follows that $C_p(\pi') \leq C_p(\pi)$, with equality if and only if $\pi \in T(\mu, \nu)$, and thus $\text{Opt}_p(\mu, \nu) \subset T(\mu, \nu)$.
2. Write $\tilde{\pi} = \iota(\pi)$. The mass $\tilde{\pi}(\{\partial\Omega\} \times \{\partial\Omega\})$ is non-negative by definition. One has for all Borel sets $A \subset \Omega$,

$$\begin{aligned} \tilde{\pi}(A \times \tilde{\Omega}) &= \tilde{\pi}(A \times \Omega) + \tilde{\pi}(A \times \{\partial\Omega\}) \\ &= \pi(A \times \Omega) + \pi(A \times \partial\Omega) = \pi(A \times \bar{\Omega}) = \mu(A) = \tilde{\mu}(A). \end{aligned}$$

Similarly, $\tilde{\pi}(\tilde{\Omega} \times B) = \tilde{\nu}(B)$ for all $B \subset \Omega$. Observe also that

$$\tilde{\pi}(\{\partial\Omega\} \times \tilde{\Omega}) = \tilde{\pi}(\{\partial\Omega\} \times \{\partial\Omega\}) + \tilde{\pi}(\{\partial\Omega\} \times \Omega) = r - \mu(\Omega) = \tilde{\mu}(\{\partial\Omega\}).$$

Similarly, $\tilde{\pi}(\tilde{\Omega} \times \{\partial\Omega\}) = \tilde{\nu}(\{\partial\Omega\})$. It gives that $\iota(\pi) \in \Pi(\tilde{\mu}, \tilde{\nu})$, so that ι is well defined. Observe that

$$\begin{aligned} \iint_{\tilde{\Omega} \times \tilde{\Omega}} d(x, y)^p d\tilde{\pi}(x, y) &= \iint_{\Omega \times \Omega} d(x, y)^p d\pi(x, y) \\ &\quad + \int_{\Omega} d(x, \partial\Omega)^p d\pi(x, \partial\Omega) \\ &\quad + \int_{\Omega} d(\partial\Omega, y)^p d\pi(\partial\Omega, y) + 0 \\ &= C_p(\pi) \text{ as } \pi \in T(\mu, \nu). \end{aligned}$$

3. Write $\pi = \kappa(\tilde{\pi})$. For $A \subset \Omega$ a Borel set,

$$\begin{aligned} \pi(A \times \bar{\Omega}) &= \pi(A \times \Omega) + \pi(A \times \partial\Omega) \\ &= \tilde{\pi}(A \times \Omega) + \tilde{\pi}(A \times \{\partial\Omega\}) = \tilde{\pi}(A \times \tilde{\Omega}) = \mu(A). \end{aligned}$$

Similarly, $\pi(\bar{\Omega} \times B) = \nu(B)$ for all $B \subset \Omega$. Therefore, $\pi \in \text{Adm}(\mu, \nu)$, and by construction, if a point $x \in \Omega$ is transported on $\partial\Omega$, it is

transported on $\mathbf{p}_{\partial\Omega}(x)$, so that $\pi \in T(\mu, \nu)$. Observe that $\mu(\Omega) + \pi(\partial\Omega \times \Omega) \leq \tilde{\pi}(\tilde{\Omega} \times \tilde{\Omega}) = r$, so that $\iota(\pi)$ is well defined. Also, $\iota(\pi) = \tilde{\pi}$, so that, according to point 2, $C_p(\pi) = \iint_{\tilde{\Omega} \times \tilde{\Omega}} d(x, y)^p d\tilde{\pi}(x, y)$.

□

Proof of Lemma 3.17. Let $\tilde{\pi} \in \Pi(\tilde{\mu}, \tilde{\nu})$. Define the set $H := \{(x, y) \in \tilde{\Omega}^2, \rho(x, y) = d(x, y)\}$, and let H^c be its complementary set in $\tilde{\Omega}^2$, i.e. the set where $\rho(x, y) = d(x, \partial\Omega) + d(\partial\Omega, y)$. Define $\tilde{\pi}' \in \mathcal{M}(\tilde{\Omega}^2)$ by, for Borel sets $A, B \subset \Omega$:

$$\begin{cases} \tilde{\pi}'(A \times B) = \tilde{\pi}((A \times B) \cap H) \\ \tilde{\pi}'(A \times \{\partial\Omega\}) = \tilde{\pi}((A \times \tilde{\Omega}) \cap H^c) + \tilde{\pi}(A \times \{\partial\Omega\}) \\ \tilde{\pi}'(\{\partial\Omega\} \times B) = \tilde{\pi}((\tilde{\Omega} \times B) \cap H^c) + \tilde{\pi}(\{\partial\Omega\} \times B). \end{cases}$$

We easily check that $\tilde{\pi}' \in \Pi(\tilde{\mu}, \tilde{\nu})$. Also, using $(a+b)^p \geq a^p + b^p$ for positive a, b , we have

$$\begin{aligned} \iint_{\tilde{\Omega} \times \tilde{\Omega}} \rho(x, y)^p d\tilde{\pi}(x, y) &= \iint_H d(x, y)^p d\tilde{\pi}(x, y) \\ &\quad + \iint_{H^c} (d(x, \partial\Omega) + d(\partial\Omega, y))^p d\tilde{\pi}(x, y) \\ &\geq \iint_H d(x, y)^p d\tilde{\pi}'(x, y) \\ &\quad + \iint_{H^c} (d(x, \partial\Omega)^p + d(y, \partial\Omega)^p) d\tilde{\pi}(x, y) \\ &= \iint_{\tilde{\Omega} \times \tilde{\Omega}} d(x, y)^p d\tilde{\pi}'(x, y) \\ &\geq \inf_{\tilde{\pi}' \in \Pi(\tilde{\mu}, \tilde{\nu})} \iint_{\tilde{\Omega} \times \tilde{\Omega}} d(x, y)^p d\tilde{\pi}'(x, y). \end{aligned}$$

We conclude by taking the infimum on $\tilde{\pi}$ that

$$W_{p,\rho}(\tilde{\mu}, \tilde{\nu}) \geq \inf_{\tilde{\pi}' \in \Pi(\tilde{\mu}, \tilde{\nu})} \iint_{\tilde{\Omega} \times \tilde{\Omega}} d(x, y)^p d\tilde{\pi}'(x, y).$$

Since $\rho(x, y) \leq d(x, y)$, it follows that

$$W_{p,\rho}^p(\tilde{\mu}, \tilde{\nu}) = \inf_{\tilde{\pi} \in \Pi(\tilde{\mu}, \tilde{\nu})} \iint_{\tilde{\Omega}^2} d(x, y)^p d\tilde{\pi}(x, y). \quad (3.21)$$

Since d is continuous, the infimum in the right hand side of (3.21) is reached [Vil08, Theorem 4.1]. Consider thus $\tilde{\pi} \in \Pi(\tilde{\mu}, \tilde{\nu})$ which realizes the infimum.

We can write, using Lemma 3.16,

$$\begin{aligned} W_{p,\rho}^p(\tilde{\mu}, \tilde{\nu}) &= \iint_{\tilde{\Omega}^2} d(x, y)^p d\tilde{\pi}(x, y) = \iint_{\bar{\Omega} \times \bar{\Omega}} d(x, y)^p d\kappa(\tilde{\pi})(x, y) \\ &\geq \inf_{\pi \in T(\mu, \nu)} \iint_{\bar{\Omega} \times \bar{\Omega}} d(x, y)^p d\pi(x, y) = \text{OT}_p^p(\mu, \nu), \end{aligned}$$

which concludes the proof. \square

Proof of Proposition 3.26. First, the convergence $\text{Pers}_\infty(\mu_n) \rightarrow \text{Pers}_\infty(\mu)$ is a consequence of the reverse triangle inequality:

$$|\text{Pers}_\infty(\mu_n) - \text{Pers}_\infty(\mu)| = |\text{OT}_\infty(\mu_n, 0) - \text{OT}_\infty(\mu, 0)| \leq \text{OT}_\infty(\mu_n, \mu),$$

which converges to 0 as n goes to ∞ .

We now prove the vague convergence. Let $f \in C_c(\Omega)$, whose support is included in some compact set K . For any $\varepsilon > 0$, there exists a L -Lipschitz function f_ε , whose support is included in K , with $\|f - f_\varepsilon\|_\infty \leq \varepsilon$. Observe that $\sup_k \mu_k(K) < \infty$ using the same arguments than for (3.14). Let $\pi_n \in \text{Opt}_\infty(\mu_n, \mu)$. We have

$$\begin{aligned} |\mu_n(f) - \mu(f)| &\leq |\mu_n(f - f_\varepsilon)| + |\mu(f - f_\varepsilon)| + |\mu_n(f_\varepsilon) - \mu(f_\varepsilon)| \\ &\leq (\mu_n(K) + \mu(K))\varepsilon + |\mu_n(f_\varepsilon) - \mu(f_\varepsilon)| \\ &\leq (\sup_k \mu_k(K) + \mu(K))\varepsilon + |\mu_n(f_\varepsilon) - \mu(f_\varepsilon)|. \end{aligned}$$

Also,

$$\begin{aligned} |\mu_n(f_\varepsilon) - \mu(f_\varepsilon)| &\leq \iint_{\bar{\Omega}^2} |f_\varepsilon(x) - f_\varepsilon(y)| d\pi_n(x, y) \\ &\leq L \iint_{(K \times \bar{\Omega}) \cup (\bar{\Omega} \times K)} d(x, y) d\pi_n(x, y) \\ &\leq LC_\infty(\pi_n)(\pi_n(K \times \bar{\Omega}) + \pi_n(\bar{\Omega} \times K)) \\ &\leq L\text{OT}_\infty(\mu_n, \mu) \left(\sup_k \mu_k(K) + \mu(K) \right) \rightarrow 0. \end{aligned}$$

This last quantity converges to 0 as $n \rightarrow \infty$ for fixed ε . Therefore, taking the lim sup in n and then letting ε go to 0, we obtain that $\mu_n(f) \rightarrow \mu(f)$. \square

Proof of Proposition 3.33. (i) \Rightarrow (ii). Let $\pi \in \text{Opt}(\mu, \nu)$. Define π' as in [FG10, Eq. (7)], that is

$$\pi' := \pi|_{\mathcal{A}} + (p_1, \mathbf{p}_{\partial\Omega} \circ p_1) \# \pi|_{\bar{\Omega} \times \bar{\Omega} \setminus \mathcal{A}} + (\mathbf{p}_{\partial\Omega} \circ p_2, p_2) \# \pi|_{\bar{\Omega} \times \bar{\Omega} \setminus \mathcal{A}},$$

where $\mathbf{p}_{\partial\Omega}$ is the orthogonal projection onto the diagonal, and p_1, p_2 are the projections of a measure onto its first and second marginal respectively. One can observe that $\pi \in \text{Adm}(\mu, \nu)$, that $\text{spt}(\pi) \subset \mathcal{A}$, and that $C_p(\pi') \leq C_p(\pi)$ with a strict inequality if $\text{spt}(\pi) \not\subset \mathcal{A}$. As such, any optimal transport plan must be supported on \mathcal{A} .

In order to prove \tilde{c} -cyclical monotonicity, we first introduce μ_n, ν_n two measures with finite mass. In this context, we know from Proposition 3.15 and Remark 3.18 that $\text{OT}_p(\mu_n, \nu_n) = W_{p, \tilde{c}}(\tilde{\mu}_n, \tilde{\nu}_n)$ and that an optimal transport plan $\tilde{\pi}_n$ between $\tilde{\mu}_n$ and $\tilde{\nu}_n$ for the cost \tilde{c} induces an optimal transport plan π_n between μ_n and ν_n for the cost c . Using [Vil08, Thm. 5.10, part (ii)], it follows that $\text{spt}(\tilde{\pi}_n)$ is \tilde{c} -cyclically monotone, and thus that $\text{spt}(\pi_n)$ is also \tilde{c} -cyclically monotone.

Now, consider $\mu, \nu \in \mathcal{M}^p$ and $\pi \in \text{Opt}_p(\mu, \nu)$. Denote by π_n the restriction of π to $(A_{1/n}^c \cup \partial\Omega) \times (A_{1/n}^c \cup \partial\Omega)$, where $A_{1/n}^c := \{x \in \Omega, d(x, \partial\Omega) > \frac{1}{n}\}$, and let μ_n, ν_n denote respectively the first and second marginal of π_n restricted to Ω . Observe that $\pi_n \in \text{Opt}_p(\mu_n, \nu_n)$, otherwise one could consider $\pi'_n \in \text{Opt}_p(\mu_n, \nu_n)$ that would satisfy $C_p(\pi'_n) < C_p(\pi_n)$ and then $\pi' = (\pi - \pi_n) + \pi'_n \in \text{Adm}(\mu, \nu)$ would achieve a strictly smaller cost than π which is supported to be optimal. Furthermore, μ_n and ν_n have finite total masses. Indeed, for any Borel subset $B \subset \Omega$, one has

$$\mu_n(B) = \pi_n(B, \overline{\Omega}) \leq \pi(B \cap A_{1/n}^c, \overline{\Omega}) = \mu(B \cap A_{1/n}^c) \leq \mu(A_{1/n}^c) < \infty,$$

as $\text{Pers}_p(\mu) < \infty$, see Remark 3.3; and a similar argument holds for ν_n . Therefore, π_n is an optimal transport plan between two finite measures and its support is thus \tilde{c} -cyclically monotone by the previous point. Eventually, $\pi_n \xrightarrow{v} \pi$ as $n \rightarrow \infty$, and thus $\text{spt}(\pi)$ is also \tilde{c} -cyclically monotone. Indeed, fix an integer $K \in \mathbb{N}$ and $(x_1, y_1) \dots (x_K, y_K) \in \text{spt}(\pi)$; for n large enough, $(x_1, y_1) \dots (x_K, y_K)$ belong to $\text{spt}(\pi_n)$, thus satisfy (3.15) (with \tilde{c} instead of c), proving the claim.

In order to prove (ii) \Rightarrow (iii), as remarked in [FG10], it suffices to observe that these statements do not deal with the mass of π and its marginals, but only about its support. As such, one can just reproduce the Step 3 of the proof of [Vil08, Thm. 5.10, part (i)] along with the computations made in the proof of [FG10, Prop. 2.3]. More precisely, it suffices to fix $(x_0, y_0) \in \text{spt}(\pi)$ and to define

$$\begin{aligned} \varphi(x) := & \sup_{K \in \mathbb{N}} \sup \{ [\tilde{c}(x_1, y_0) - \tilde{c}(x_0, y_0)] + [\tilde{c}(x_2, y_1) - \tilde{c}(x_1, y_1)] \\ & + \dots + \tilde{c}(x, y_K) - \tilde{c}(x_K, y_K) \}, \end{aligned}$$

where the second supremum is taken over $(x_1, y_1) \dots (x_K, y_K) \in \text{spt}(\pi)$. Observe that φ is \tilde{c} -concave and that $\text{spt}(\pi) \subset \partial^{\tilde{c}}\varphi$. Furthermore, for $x, y \in \partial\Omega$, one has $\varphi(x) + \varphi^{\tilde{c}}(y) = \tilde{c}(x, y) = 0$. As such, up to adding a constant, we have that φ and $\varphi^{\tilde{c}}$ both vanish on $\partial\Omega$. To prove that these observations remain true when replacing \tilde{c} by c , following [FG10], we observe that $\text{spt}(\pi) \subset \partial^{\tilde{c}}\varphi \cap \mathcal{A} \subset \partial^c\varphi$. Then, recalling that $\varphi(x) = 0$ for $x \in \partial\Omega$, it follows $\varphi^c(x) = \varphi(x) + \varphi^c(x) = c(x, x) = 0$, proving that φ^c also vanishes on $\partial\Omega$.

Finally, to prove (iii) \Rightarrow (i), consider $\pi, \pi' \in \text{Adm}(\mu, \nu)$ (supported on E_Ω) and assume that $\pi \subset \partial^c\varphi$ for some c -concave function φ with φ, φ^c that vanish on $\partial\Omega$. Observe that for all x, y ,

$$\varphi(x) \leq c(x, \partial\Omega), \quad \varphi^c(y) \leq c(\partial\Omega, y),$$

so that $\mu(\varphi), \mu(\varphi^c) < \infty$ for any $\mu \in \mathcal{M}^p$. Finally, observe that for any $x, y \in \bar{\Omega}$, $\varphi(x) + \varphi^c(y) \leq c(x, y)$ (while equality holds on $\partial^c\varphi$). It allows us to write

$$\begin{aligned} C_p(\pi) &= \iint_{\partial^c\varphi} c(x, y) d\pi(x, y) \\ &= \iint_{\Omega \times \bar{\Omega}} \varphi(x) d\pi(x, y) + \iint_{\bar{\Omega} \times \Omega} \varphi^c(y) d\pi(x, y) \\ &= \int \varphi(x) d\mu(x) + \int \varphi^c(y) d\nu(y) \\ &= \iint_{\Omega \times \bar{\Omega}} \varphi(x) d\pi'(x, y) + \iint_{\bar{\Omega} \times \Omega} \varphi^c(y) d\pi'(x, y) \\ &= \iint_{\bar{\Omega} \times \bar{\Omega}} (\varphi(x) + \varphi^c(y)) d\pi'(x, y) \\ &\leq \iint c(x, y) d\pi'(x, y) = C_p(\pi'). \end{aligned}$$

Therefore, $\pi \in \text{Opt}(\mu, \nu)$.

The last assertion is known as *stability of optimal transport* in standard optimal transport literature [Vil08, Thm. 5.20], and mentioned in [FG10]. It is an immediate corollary of the equivalence between (i) and (ii): let $\pi_n \in \text{Opt}(\mu_n, \nu_n)$. As such, for all n , the support of π_n is a subset of \mathcal{A} and is \tilde{c} -cyclically monotone. Up to extracting a subsequence, $(\pi_n)_n$ converges vaguely to some π , which also has its support included in \mathcal{A} and \tilde{c} -cyclically monotone, as \tilde{c} is continuous. Therefore, $\pi \in \text{Opt}(\mu, \nu)$. \square

CHAPTER 4

FRÉCHET MEANS IN THE SPACE OF PERSISTENCE MEASURES

Abstract

In this chapter, we state the existence of Fréchet means for probability distributions supported on \mathcal{M}^p . We start with the finite case (i.e. averaging finitely many persistence measures) and then extend the result to any probability distribution with finite p -th moment. We then study the specific case of distribution supported on \mathcal{D}^p (i.e. averaging persistence diagrams) and show that in the finite setting, the set of Fréchet means is a convex set whose extreme points are in \mathcal{D}^d (i.e. are actual persistence diagrams). As in the previous chapter, some technical proofs have been deferred to a specific section (§4.4). Note that computational aspects are not addressed here: they will be covered in the next part of the manuscript, which is dedicated to applications.

REMARK 4.1. *In this chapter, we will assume that $1 < p < \infty$ and $1 < q < \infty$ (recall $d(\cdot, \cdot) = \|\cdot - \cdot\|_q$). These assumptions will ensure that (i) the projection of $x \in \Omega$ onto $\partial\Omega$ is uniquely defined and (ii) the Fréchet mean of k points $x_1 \dots x_k$ in $\bar{\Omega}$, i.e. minimizer of $x \mapsto \sum_{i=1}^k \|x - x_i\|_q^p$, is also uniquely defined; two facts used in our proofs.*

Recall that $(\mathcal{M}^p, \text{OT}_p)$ is a Polish space, and let W_p denote the Wasserstein distance (see Section 2.2.2) between probability measures supported on $(\mathcal{M}^p, \text{OT}_p)$. We denote by $\mathcal{W}^p(\mathcal{M}^p)$ the space of probability measures P supported on \mathcal{M}^p , equipped with the W_p metric, which are at a finite

distance from $\delta_{\mathbf{0}}$ —the Dirac mass supported on the empty diagram—i.e.

$$W_p^p(P, \delta_{\mathbf{0}}) = \int_{\nu \in \mathcal{M}^p} \text{OT}_p^p(\nu, \mathbf{0}) dP(\nu) = \int_{\nu \in \mathcal{M}^p} \text{Pers}_p(\nu) dP(\nu) < \infty.$$

DEFINITION 4.2. *Consider $P \in \mathcal{W}^p(\mathcal{M}^p)$. A measure $\mu^* \in \mathcal{M}^p$ is a Fréchet mean of P if it minimizes the energy functional*

$$\mathcal{E} : \mu \in \mathcal{M}^p \mapsto \int_{\nu \in \mathcal{M}^p} \text{OT}_p^p(\mu, \nu) dP(\nu).$$

4.1 Fréchet means in the finite case

Let P be of the form $\sum_{i=1}^N \lambda_i \delta_{\mu_i}$ with $N \in \mathbb{N}$, μ_i a persistence measure of finite mass m_i , and $(\lambda_i)_i$ non-negative weights that sum to 1. Let \mathcal{E} be the corresponding energy functional. Define $m_{\text{tot}} := \sum_{i=1}^N m_i$. To prove the existence of Fréchet means for such a P , we show that, in this case, Fréchet means correspond to Fréchet means for the Wasserstein distance of some distribution on $\mathcal{M}_{m_{\text{tot}}}^p(\tilde{\Omega})$, the set of measures on $\tilde{\Omega}$ that all have the same mass m_{tot} (see Section 3.2), a problem well studied in the literature [AC11, CE10, COO15a]. Recall from Section 3.2 that $\tilde{\Omega}$ denotes the closed half-plane $\bar{\Omega}$ where the points of the boundary $\partial\Omega$ have been identified. This space is naturally equipped with a metric ρ deduced from d by enforcing the triangle inequality, that is

$$\rho(x, y) = \min(d(x, y), d(x, \mathbf{p}_{\partial\Omega}(x)) + d(\mathbf{p}_{\partial\Omega}(y), y)).$$

We start with a lemma which claims that if a measure μ has too much mass (larger than m_{tot}), then it cannot be a Fréchet mean of $\mu_1 \dots \mu_N$. Its proof can be found in Section 4.4. The idea of the proof is to show that if a measure μ has some mass that is mapped to the diagonal in each transport plan between μ and μ_i , then we can build a measure μ' by “removing” this mass, and then observe that such a measure μ' has a smaller energy $\mathcal{E}(\mu')$.

LEMMA 4.3. *We have $\inf\{\mathcal{E}(\mu), \mu \in \mathcal{M}^p\} = \inf\{\mathcal{E}(\mu), \mu \in \mathcal{M}_{\leq m_{\text{tot}}}^p\}$.*

Let $W_{p,\rho}$ denote the Wasserstein distance between measures with same mass supported on the metric space $(\tilde{\Omega}, \rho)$ (see Sections 2.2.2 and 3.2).

PROPOSITION 4.4. *Let*

$$\begin{aligned} \Psi : \mathcal{M}_{\leq m_{\text{tot}}}^p &\rightarrow \mathcal{M}_{m_{\text{tot}}}^p(\tilde{\Omega}) \\ \mu &\mapsto \mu + (m_{\text{tot}} - \mu(\tilde{\Omega}))\delta_{\partial\Omega}. \end{aligned}$$

The functionals

$$\mathcal{E} : \mu \in \mathcal{M}_{\leq m_{\text{tot}}}^p \mapsto \sum_{i=1}^N \lambda_i \text{OT}_p^p(\mu, \mu_i)$$

and

$$\mathcal{F} : \tilde{\mu} \in \mathcal{M}_{m_{\text{tot}}}^p(\tilde{\Omega}) \mapsto \sum_{i=1}^N \lambda_i W_{p,\rho}^p(\tilde{\mu}, \Psi(\mu_i)),$$

have the same infimum values and $\arg \min \mathcal{E} = \Psi^{-1}(\arg \min \mathcal{F})$.

The existence of minimizers $\tilde{\mu}$ of \mathcal{F} , that is “Wasserstein barycenter” (i.e. Fréchet means for the Wasserstein distance) of $\hat{P} := \sum_{i=1}^N \lambda_i \delta_{\tilde{\mu}_i}$, is well-known (see [AC11, Theorem 8]). Proposition 4.4 asserts that $\Psi^{-1}(\tilde{\mu})$ is a minimizer of \mathcal{E} on $\mathcal{M}_{\leq m_{\text{tot}}}^p$, and thus a Fréchet mean of P according to Lemma 4.3. We therefore have the following corollary, namely the existence of Fréchet means in the finite case.

COROLLARY 4.5. *Let $\mu_1 \dots \mu_N$ be N persistence measures in \mathcal{M}^p , and let $\lambda_1 \dots \lambda_N$ be non-negative weights that sum to one. The distribution $P = \sum_{i=1}^N \lambda_i \delta_{\mu_i} \in \mathcal{W}^p(\mathcal{M}^p)$ admits a Fréchet mean.*

4.2 Existence and consistency of Fréchet means

We now extend the results of the previous section to the Fréchet means of general probability distributions supported on \mathcal{M}^p . First, we show a consistency result, in the vein of [LGL16, Theorem 3].

PROPOSITION 4.6. *Let $(P_n)_n, P$ be probability measures in $\mathcal{W}^p(\mathcal{M}^p)$. Assume that each P_n has a Fréchet mean μ_n and that $W_p(P_n, P) \rightarrow 0$. Then, the sequence $(\mu_n)_n$ is relatively compact in $(\mathcal{M}^p, \text{OT}_p)$, and any limit of a converging subsequence is a Fréchet mean of P .*

To prove this result, we introduce the following lemma, whose proof can be found in Section 4.4.

LEMMA 4.7. *Let $\mu, \mu_1, \mu_2, \dots \in \mathcal{M}^p$. Then, $\text{OT}_p(\mu_n, \mu) \rightarrow 0$ if and only if $\mu_n \xrightarrow{v} \mu$ and there exists a persistence measure $\nu \in \mathcal{M}^p$ such that $\text{OT}_p(\mu_n, \nu) \rightarrow \text{OT}_p(\mu, \nu)$.*

Proof of Proposition 4.6. We first prove the relative compactness of $(\mu_n)_n$. Following Proposition B.4, we can show that for every compact $K \subset \Omega$, $\sup\{\mu_n(K)\} < \infty$. Consider a compact set $K \subset \Omega$. We have, because of (3.6),

$$\begin{aligned} \mu_n(K)^{\frac{1}{p}} &\leq \frac{1}{d(K, \partial\Omega)} \text{OT}_p(\mu_n, \mathbf{0}) = \frac{1}{d(K, \partial\Omega)} W_p(\delta_{\mu_n}, \delta_{\mathbf{0}}) \\ &\leq \frac{1}{d(K, \partial\Omega)} (W_p(\delta_{\mu_n}, P_n) + W_p(P_n, \delta_{\mathbf{0}})). \end{aligned}$$

Since μ_n is a Fréchet mean of P_n , it minimizes $\{W_p(\delta_\nu, P_n), \nu \in \mathcal{M}^p\}$, and in particular $W_p(\delta_{\mu_n}, P_n) \leq W_p(\delta_{\mathbf{0}}, P_n)$. Furthermore, since we assume that $W_p(P_n, P) \rightarrow 0$, we have in particular that $\sup_n W_p(P_n, \delta_{\mathbf{0}}) < \infty$. As a consequence $\sup_n \mu_n(K) < \infty$, and it follows (Proposition B.4) that the sequence $(\mu_n)_n$ is relatively compact for the vague topology.

Without relabeling, assume thus that $(\mu_n)_n \xrightarrow{v} \mu$. One must show that μ is a Fréchet mean of P . At this stage, we can follow the same computations as [LGL16, Proof of Theorem 3]. Let ν be a persistence measure. We write

$$\begin{aligned} \int \text{OT}_p^p(\nu, \xi) dP(\xi) &= W_p^p(\delta_\nu, P) \\ &= \lim_{n \rightarrow \infty} W_p^p(\delta_\nu, P_n) \\ &\geq \lim_{n \rightarrow \infty} W_p^p(\delta_{\mu_n}, P_n) \\ &= \lim_{n \rightarrow \infty} \int \text{OT}_p^p(\mu_n, \xi) dP_n(\xi) \\ &\geq \int \liminf_{n \rightarrow \infty} \text{OT}_p^p(\mu_n, \xi) dP_n(\xi) \\ &\geq \int \text{OT}_p^p(\mu, \xi) dP(\xi) \quad \text{by lower semi-continuity.} \end{aligned}$$

This proves that μ must be a barycenter of P . Furthermore, taking $\nu = \mu$ in the above series of inequalities yields

$$W_p(\delta_{\mu_n}, P_n) \rightarrow W_p(\delta_\mu, P),$$

and then by triangle inequality, and using that $W_p(P_n, P) \rightarrow 0$,

$$W_p(\delta_{\mu_n}, P) \rightarrow W_p(\delta_\mu, P).$$

Therefore, we can write

$$\begin{aligned}
\int \text{OT}_p^p(\mu, \xi) dP(\xi) &= W_p^p(\delta_\mu, P) \\
&= \lim_{n \rightarrow \infty} W_p^p(\delta_{\mu_n}, P) \\
&= \lim_{n \rightarrow \infty} \int \text{OT}_p^p(\mu_n, \xi) dP(\xi) \\
&\geq \int \liminf_{n \rightarrow \infty} \text{OT}_p^p(\mu_n, \xi) dP(\xi) \\
&\geq \int \text{OT}_p^p(\mu, \xi) dP(\xi),
\end{aligned}$$

forcing in particular that

$$\int \underbrace{\liminf_{n \rightarrow \infty} \text{OT}_p^p(\mu_n, \xi) - \text{OT}_p^p(\mu, \xi)}_{\geq 0} dP(\xi) = 0,$$

and thus, P -almost surely, we have

$$\liminf_{n \rightarrow \infty} \text{OT}_p^p(\mu_n, \xi) = \text{OT}_p^p(\mu, \xi).$$

We eventually obtain the existence of some $\nu \in \mathcal{M}^p$ such that, up to a subsequence,

$$\text{OT}_p(\mu_n, \nu) \rightarrow \text{OT}_p(\mu, \nu).$$

Lemma 4.7 gives the desired conclusion. \square

Proposition 4.6 allows us to generalize existence of barycenters in the finite case (Section 4.1) to general probability distributions supported on \mathcal{M}^p .

THEOREM 4.8. *For any probability distribution P supported on \mathcal{M}^p with finite p -th moment, the set of Fréchet means of P is a non-empty compact convex set of \mathcal{M}^p .*

Note that here, *convexity* must be understood in the sense of linear combination of measures, that is a convex set $C \subset \mathcal{M}^p$ is such that for any two measures $\mu, \nu \in C$, and any $t \in (0, 1)$, one has $(1-t)\mu + t\nu \in C$.

Proof. We first prove the non-emptiness. Let $P = \sum_{i=1}^N \lambda_i \mu_i$ be a probability measure on \mathcal{M}^p with finite support μ_1, \dots, μ_N . According to Proposition 3.14, there exists sequences $(\mu_i^{(n)})_n$ in \mathcal{M}_f^p with $\text{OT}_p(\mu_i^{(n)}, \mu_i) \rightarrow 0$.

As a consequence of the result of Section 4.1, the probability measures $P^{(n)} := \sum_i \lambda_i \delta_{\mu_i^{(n)}}$ admit Fréchet means. Furthermore, $W_p^p(P^{(n)}, P) \leq \sum_i \lambda_i \text{OT}_p^p(\mu_i^{(n)}, \mu_i)$ so that this quantity converges to 0 as $n \rightarrow \infty$. It follows from Proposition 4.6 that P admits a Fréchet mean.

If P has infinite support, following [LGL16], it can be approximated (in W_p) by an empirical probability measure $P_n = \frac{1}{n} \sum_{i=1}^n \delta_{\mu_i}$ where the μ_i are i.i.d. from P . We know that P_n admits a Fréchet mean since its support is finite, and thus, applying Proposition 4.6 once again, we obtain that P admits a Fréchet mean.

Finally, the compactness of the set of Fréchet means follows from Proposition 4.6 and the convexity of the set from the convexity of OT_p^p (Lemma 3.13). \square

4.3 Fréchet means in \mathcal{D}^p

We now prove the existence of Fréchet means for distributions of persistence diagram (i.e. probability distributions supported on \mathcal{D}^p), extending the results of [MMH11], in which authors prove their existence for specific probability distributions (namely distributions with compact support or specific rates of decay). Theorem 4.10 below asserts two different things: that $\arg \min\{\mathcal{E}(a), a \in \mathcal{D}^p\}$ is non empty, and that

$$\min\{\mathcal{E}(a), a \in \mathcal{D}^p\} = \min\{\mathcal{E}(\mu), \mu \in \mathcal{M}^p\},$$

that is an optimal persistence measure cannot perform strictly better than an optimal persistence diagram when averaging persistence diagrams. As for Fréchet means in \mathcal{M}^p , we start with the finite case. The following lemma actually gives a geometric description of the set of Fréchet means obtained when averaging a finite number of finite diagrams.

LEMMA 4.9. *Consider $a_1, \dots, a_N \in \mathcal{D}_f$, weights $(\lambda_i)_i$ that sum to 1, and let $P := \sum_{i=1}^N \lambda_i \delta_{a_i}$. Then, the set of minimizers of $\mu \mapsto \sum_{i=1}^N \lambda_i \text{OT}_p^p(\mu, a_i)$ is a non empty convex subset of \mathcal{M}_f whose extreme points belong to \mathcal{D}_f . In particular, P admits a Fréchet mean in \mathcal{D}_f .*

The proof of this lemma relies on *integer programming* results and is found in Section 4.4. Note that, as a straightforward consequence, if P has a unique minimizer in \mathcal{D}_f (which is generically true [Tur13]), then so it does in \mathcal{M}_f (and this unique minimizer is a persistence diagram).

THEOREM 4.10. *For any probability distribution P supported on \mathcal{D}^p with finite p -th moment, the set of Fréchet means of P contains an element of \mathcal{D}^p . Furthermore, if P is supported on a finite set of finite persistence diagrams, then the set of the Fréchet means of P is a convex set whose extreme points are in \mathcal{D}^p .*

Proof. The second assertion of the theorem is stated in Lemma 4.9. To prove the existence of a Fréchet mean which is a persistence diagram, we argue as in the proof of Theorem 4.8, using additionally the fact that \mathcal{D}^p is closed in \mathcal{M}^p (Corollary 3.9). \square

4.4 Proofs

Proof of Lemma 4.3. Let $\mu \in \mathcal{M}^p$. Let $\pi_i \in \text{Opt}_p(\mu_i, \mu)$ for $i = 1, \dots, N$. The measure $A \subset \Omega \mapsto \pi_i(\partial\Omega \times A)$ is absolutely continuous with respect to μ . Therefore, it has a density f_i with respect to μ . Define for $A \subset \Omega$ a Borel set,

$$\mu'(A) := \mu(A) - \int_A \min_j f_j(x) d\mu(x),$$

and, for $i = 1, \dots, N$, a measure π'_i , equal to π_i on $\Omega \times \bar{\Omega}$ and which satisfies for $A \subset \Omega$ a Borel set,

$$\pi'_i(\partial\Omega \times A) = \pi'_i(\mathbf{p}_{\partial\Omega}(A) \times A) := \pi_i(\partial\Omega \times A) - \int_A \min_j f_j(x) d\mu(x),$$

where s is the orthogonal projection on $\partial\Omega$. As $\pi_i(\partial\Omega \times A) = \int_A f_i(x) d\mu(x)$, $\pi'_i(A)$ is non-negative, and as $\pi_i(\partial\Omega \times A) \leq \mu(A)$, it follows that $\mu'(A)$ is non-negative. To prove that $\pi'_i \in \text{Adm}(\mu_i, \mu')$, it is enough to check that for $A \subset \Omega$, $\pi'_i(\bar{\Omega} \times A) = \mu'(A)$:

$$\begin{aligned} \pi'_i(\bar{\Omega} \times A) &= \pi_i(\Omega \times A) + \pi_i(\partial\Omega \times A) - \int_A \min_j f_j(x) d\mu(x) \\ &= \mu(A) - \int_A \min_j f_j(x) d\mu(x) = \mu'(A). \end{aligned}$$

Also,

$$\begin{aligned}
\mu'(\Omega) &= \int_{\Omega} (1 - \min_j f_j) d\mu(x) \leq \sum_{j=1}^N \int_{\Omega} (1 - f_j) d\mu(x) \\
&= \sum_{j=1}^N (\mu(\Omega) - \pi_j(\partial\Omega \times \Omega)) = \sum_{j=1}^N (\pi_j(\bar{\Omega} \times \Omega) - \pi_j(\partial\Omega \times \Omega)) \\
&= \sum_{j=1}^N \pi_j(\Omega \times \Omega) \leq \sum_{j=1}^N \pi_j(\Omega \times \bar{\Omega}) = \sum_{j=1}^N m_j = m_{\text{tot}}.
\end{aligned}$$

and thus $\mu'(\Omega) \leq m_{\text{tot}}$. To conclude, observe that

$$\begin{aligned}
\mathcal{E}(\mu') &\leq \sum_{i=1}^N \lambda_i C_p(\pi'_i) = \sum_{i=1}^N \lambda_i \left(\iint_{\Omega \times \bar{\Omega}} d(x, y)^p d\pi_i(x, y) \right. \\
&\quad \left. + \iint_{\partial\Omega \times \Omega} d(x, y)^p d\pi_i(x, y) - \int_{\Omega} d(x, \partial\Omega)^p \min_j f_j(x) d\mu(x) \right) \\
&\leq \sum_{i=1}^N \lambda_i C_p(\pi) = \mathcal{E}(\mu).
\end{aligned}$$

□

Proof of Proposition 4.4. Let G be the set of $\mu \in \mathcal{M}^p$ such that, for all i , there exists $\pi_i \in \text{Opt}_p(\mu_i, \mu)$ with $\pi_i(\Omega, \partial\Omega) = 0$. By point 2 of Lemma 3.16, for $\mu \in G$ and $\pi_i \in \text{Opt}_p(\mu_i, \mu)$ with $\pi_i(\Omega, \partial\Omega) = 0$, $\iota(\pi_i)$ is well defined and satisfies

$$\text{OT}_p^p(\mu_i, \mu) = C_p(\pi_i) = \iint_{\bar{\Omega} \times \bar{\Omega}} d(x, y)^p d\iota(\pi_i)(x, y) \geq \tilde{C}_p(\iota(\pi_i)) \geq W_{p,\rho}^p(\tilde{\mu}_i, \tilde{\mu}),$$

so that $\mathcal{F}(\Psi(\mu)) \leq \mathcal{E}(\mu)$. As, by Lemma 3.17, $\mathcal{E} \leq \mathcal{F} \circ \Psi$, we therefore have $\mathcal{E}(\mu) = \mathcal{F}(\Psi(\mu))$ for $\mu \in G$.

We now show that if $\mu \notin G$, then there exists $\mu' \in \mathcal{M}^p$ with $\mathcal{E}(\mu') < \mathcal{E}(\mu)$. Let $\mu \notin G$ and $\pi_i \in \text{Opt}_p(\mu_i, \mu)$. Assume that for some i , we have $\pi_i(\Omega, \partial\Omega) > 0$, and introduce $\nu \in \mathcal{M}^p$ defined as $\nu(A) = \pi_i(A, \partial\Omega)$ for $A \subset \Omega$. Define

$$T_i : \Omega \ni x \mapsto \arg \min_{y \in \Omega} \left\{ \lambda_i d(x, y)^p + \sum_{j \neq i} \lambda_j d(y, \partial\Omega)^p \right\} \in \Omega.$$

Note that since $p > 1$, this map is well defined (the minimizer is unique due to strict convexity) and continuous thus measurable. Consider the measure




Figure 4.1: Global picture of the proof. The main idea is to observe that the cost induced by π_i (red) is strictly larger than the sum of costs induced by the π'_i s (blue), which leads to a strictly better energy.

$\mu' = \mu + (T_{i\#}\nu)$, where $T_{i\#}\nu$ is the push-forward of ν by the map T_i . Consider the transport plan π'_i deduced from π_i where ν is transported onto $T_{i\#}\nu$ instead of being transported to $\partial\Omega$ (see Figure 4.1). More precisely, π'_i is the measure on $\bar{\Omega} \times \bar{\Omega}$ defined by, for Borel sets $A, B \subset \Omega$:

$$\begin{aligned}\pi'_i(A \times B) &= \pi_i(A \times B) + \nu(A \cap T_i^{-1}(B)), \\ \pi'_i(A \times \partial\Omega) &= 0, \quad \pi'_i(\partial\Omega \times B) = \pi_i(\partial\Omega \times B).\end{aligned}$$

We have $\pi'_i \in \text{Adm}(\mu_i, \mu')$. Indeed, for Borel sets $A, B \subset \Omega$:

$$\pi'_i(A \times \bar{\Omega}) = \pi'_i(A \times \Omega) = \pi_i(A \times \Omega) + \nu(A) = \pi_i(A \times \bar{\Omega}) = \mu_i(A),$$

and

$$\begin{aligned}\pi'_i(\bar{\Omega} \times B) &= \pi'_i(\Omega \times B) + \pi'_i(\partial\Omega \times B) \\ &= \pi_i(\Omega \times B) + \nu(T_i^{-1}(B)) + \pi_i(\partial\Omega \times B) \\ &= \mu(B) + T_{i\#}\nu(B) = \mu'(B).\end{aligned}$$

Using π'_i instead of π_i changes the transport cost by the quantity

$$\int_{\Omega} [d(x, T_i(x))^p - d(x, \partial\Omega)^p] d\nu(x) < 0.$$

In a similar way, we define for $j \neq i$ the plan $\pi'_j \in \text{Adm}(\mu_j, \mu')$ by transporting the mass induced by the newly added $(T_{i\#}\nu)$ to the diagonal

$\partial\Omega$. Using these modified transport plans increases the total cost by

$$\sum_{j \neq i} \lambda_j \int_{\Omega} d(T_j(x), \partial\Omega)^p d\nu(x).$$

One can observe that

$$\int_{\Omega} \left[\lambda_i (d(x, T_i(x))^p - d(x, \partial\Omega)^p) + \sum_{j \neq i} \lambda_j d(T_j(x), \partial\Omega)^p \right] d\nu(x) < 0$$

due to the definition of T and $\nu(\Omega) > 0$.

Therefore, the total transport cost induced by the $(\pi'_i)_{i=1 \dots N}$ is strictly less than $\mathcal{E}(\mu)$, and thus $\mathcal{E}(\mu') < \mathcal{E}(\mu)$. Finally, we have

$$\inf_{\mu \in \mathcal{M}_{\leq m_{tot}}^p} \mathcal{E}(\mu) = \inf_{\mu \in G} \mathcal{E}(\mu) = \inf_{\mu \in G} \mathcal{F}(\Psi(\mu)) \geq \inf_{\mu \in \mathcal{M}_{\leq m_{tot}}^p} \mathcal{F}(\Psi(\mu)) \geq \inf_{\mu \in \mathcal{M}_{\leq m_{tot}}^p} \mathcal{E}(\mu),$$

where the last inequality comes from $\mathcal{F} \circ \Psi \geq \mathcal{E}$ (Lemma 3.17). Therefore, $\inf \mathcal{E} = \inf \mathcal{F} \circ \Psi$, which is equal to $\inf \mathcal{F}$, as Ψ is a bijection. Also, if μ is a minimizer of \mathcal{E} (should it exist), then $\mu \in G$ and $\mathcal{E}(\mu) = \mathcal{F}(\Psi(\mu))$. Therefore, as the infimum are equal, $\Psi(\mu)$ is a minimizer of \mathcal{F} . Reciprocally, if $\tilde{\mu}$ is a minimizer of \mathcal{F} , then, by Lemma 3.17, $\mathcal{F}(\tilde{\mu}) \geq \mathcal{E}(\Psi^{-1}(\tilde{\mu}))$, and, as the infimum are equal, $\Psi^{-1}(\tilde{\mu})$ is a minimizer of \mathcal{E} . \square




Figure 4.2: Partition of Ω used in the proof of Lemma 4.7.

Proof of Lemma 4.7. For the direct implication, take $\nu = \mathbf{0}$ and apply Theorem 3.7.

Let us prove the converse implication. Assume that $\mu_n \xrightarrow{v} \mu$ and $\text{OT}_p(\mu_n, \nu) \rightarrow \text{OT}_p(\mu, \nu)$ for some $\nu \in \mathcal{D}^p$. The vague convergence of $(\mu_n)_n$ implies that $\mu^{[p]}$ is the only possible accumulation point for weak convergence of the sequence $(\mu_n^{[p]})_n$. Therefore, it is sufficient to show that the sequence $(\mu_n^{[p]})_n$ is relatively compact for weak convergence (i.e. tight and bounded in total variation, see Proposition B.5). Indeed, this would mean that $(\mu_n^{[p]})_n$ converges weakly to $\mu^{[p]}$, or equivalently by Proposition B.6 that $\mu_n \xrightarrow{v} \mu$ and $\text{Pers}_p(\mu_n) \rightarrow \text{Pers}_p(\mu)$. The conclusion is then obtained thanks to Theorem 3.7.

Thus, let $(\mu_n)_n$ be any subsequence and $(\pi_n)_n$ be corresponding optimal transport plans between μ_n and ν . The vague convergence of $(\mu_n)_n$ implies that $(\pi_n)_n$ is relatively compact with respect to the vague topology on E_Ω . Let π be a limit of any converging subsequence of $(\pi_n)_n$, which indexes are still denoted by n . One can prove that $\pi \in \text{Opt}(\mu, \nu)$ (see [FG10, Prop. 2.3]). For $r > 0$, define $A_r := \{x \in \Omega, d(x, \partial\Omega) \leq r\}$ and write \bar{A}_r for $A_r \cup \partial\Omega$. Consider $\eta > 1$. We can write

$$\begin{aligned}
& \int_{A_r} d(x, \partial\Omega)^p d\mu_n(x) = \iint_{A_r \times \bar{\Omega}} d(x, \partial\Omega)^p d\pi_n(x, y) \\
&= \iint_{A_r \times (\Omega \setminus A_{\eta r})} d(x, \partial\Omega)^p d\pi_n(x, y) + \iint_{\bar{A}_r \times \bar{A}_{\eta r}} d(x, \partial\Omega)^p d\pi_n(x, y) \\
&\stackrel{(*)}{\leq} \frac{1}{(\eta - 1)^p} \iint_{A_r \times (\Omega \setminus A_{\eta r})} d(x, y)^p d\pi_n(x, y) + \iint_{\bar{A}_r \times \bar{A}_{\eta r}} d(x, \partial\Omega)^p d\pi_n(x, y) \\
&\leq \frac{1}{(\eta - 1)^p} \text{OT}_p^p(\mu_n, \nu) \\
&\quad + 2^{p-1} \left(\iint_{\bar{A}_r \times \bar{A}_{\eta r}} d(x, y)^p d\pi_n(x, y) + \iint_{\bar{A}_r \times \bar{A}_{\eta r}} d(y, \partial\Omega)^p d\pi_n(x, y) \right) \\
&\leq \frac{1}{(\eta - 1)^p} \text{OT}_p^p(\mu_n, \nu) \\
&\quad + 2^{p-1} \left(\text{OT}_p^p(\mu_n, \nu) - \iint_{E_\Omega \setminus (\bar{A}_r \times \bar{A}_{\eta r})} d(x, y)^p d\pi_n(x, y) + \int_{A_{\eta r}} d(y, \partial\Omega)^p d\nu(y) \right)
\end{aligned}$$

where $(*)$ holds because $d(x, y) \geq (\eta - 1)r \geq (\eta - 1)d(x, \partial\Omega)$ for $(x, y) \in$

$A_r \times A_{\eta r}^c$. Therefore,

$$\begin{aligned} \limsup_{n \rightarrow \infty} \int_{A_r} d(x, \partial\Omega)^p d\mu_n(x) &\leq \frac{1}{(\eta - 1)^p} \text{OT}_p^p(\mu, \nu) + 2^{p-1} \left(\text{OT}_p^p(\mu, \nu) \right. \\ &\quad - \iint_{E_\Omega \setminus (\bar{A}_r \times \bar{A}_{\eta r})} d(x, y)^p d\pi(x, y) \\ &\quad \left. + \int_{A_{\eta r}} d(y, \partial\Omega)^p d\nu(y) \right) \end{aligned}$$

Note that at the last line, we used the Portmanteau theorem (see Proposition B.7) on the sequence of measures $(d(x, y)^p d\pi_n(x, y))_n$ for the open set $E_\Omega \setminus (\bar{A}_r \times \bar{A}_{\eta r})$. Letting $r \rightarrow 0$, then η goes to infinity, one obtains

$$\lim_{r \rightarrow 0} \limsup_{n \rightarrow \infty} \int_{A_r} d(x, \partial\Omega)^p d\mu_n(x) = 0.$$

The second part consists in showing that there can not be mass escaping “at infinity” in the subsequence $(\mu_n^{[p]})_n$. Fix $r, M > 0$. For $x \in \Omega$, denote $\mathbf{p}_{\partial\Omega}(x)$ the projection of x on $\partial\Omega$. Pose

$$K_{M,r} := \{x \in \Omega \setminus A_r, d(x, \partial\Omega) < M, d(\mathbf{p}_{\partial\Omega}(x), 0) < M\}$$

and $L_{M,r}$ the closure of $\Omega \setminus (A_r \cup K_{M,r})$ (see Figure 4.2). For $r' > 0$,

$$\begin{aligned} \int_{L_{M,r}} d(x, \partial\Omega)^p d\mu_n(x) &= \iint_{L_{M,r} \times \bar{\Omega}} d(x, \partial\Omega)^p d\pi_n(x, y) \\ &= \iint_{L_{M,r} \times (L_{M/2,r'} \cup \bar{A}_{r'})} d(x, \partial\Omega)^p d\pi_n(x, y) + \iint_{L_{M,r} \times K_{M/2,r'}} d(x, \partial\Omega)^p d\pi_n(x, y) \\ &\leq 2^{p-1} \iint_{L_{M,r} \times (L_{M/2,r'} \cup \bar{A}_{r'})} d(x, y)^p d\pi_n(x, y) \\ &\quad + 2^{p-1} \iint_{L_{M,r} \times (L_{M/2,r'} \cup \bar{A}_{r'})} d(\partial\Omega, y)^p d\pi_n(x, y) \\ &\quad + \iint_{L_{M,r} \times K_{M/2,r'}} d(x, \partial\Omega)^p d\pi_n(x, y). \end{aligned}$$

We treat the three parts of the sum separately. As before, taking the lim sup in n and letting $M \rightarrow \infty$, the first part of the sum converges to 0 (apply

the Portmanteau theorem on the open set $E_\Omega \setminus (L_{M,r} \times (L_{M/2,r'} \cup \overline{A_{r'}}))$. The second part is less than or equal to

$$2^{p-1} \int_{L_{M/2,r'} \cup A_{r'}} d(y, \partial\Omega)^p d\nu(y),$$

which converges to 0 as $M \rightarrow \infty$ and $r' \rightarrow 0$. For the third part, notice that if $(x, y) \in L_{M,r} \times K_{M/2,r'}$, then

$$d(x, \partial\Omega) \leq d(x, \mathbf{p}_{\partial\Omega}(y)) \leq d(x, y) + d(y, \mathbf{p}_{\partial\Omega}(y)) \leq d(x, y) + \frac{M}{2} \leq 2d(x, y).$$

Therefore,

$$\begin{aligned} \iint_{L_{M,r} \times K_{M/2,r'}} d(x, \partial\Omega)^p d\pi_n(x, y) &\leq 2^p \iint_{L_{M,r} \times K_{M/2,r'}} d(x, y)^p d\pi_n(x, y) \\ &\leq 2^p \iint_{L_{M,r} \times \overline{\Omega}} d(x, y)^p d\pi_n(x, y). \end{aligned}$$

As before, it is shown that $\limsup_n \iint_{L_{M,r} \times \overline{\Omega}} d(x, y)^p d\pi_n(x, y)$ converges to 0 when M goes to infinity by applying the Portmanteau theorem on the open set $E_\Omega \setminus (L_{M,r} \times \overline{\Omega})$.

Finally, we have shown that by taking r small enough and M large enough, one can find a compact set $\overline{K_{M,r}}$ such that $\int_{\Omega \setminus \overline{K_{M,r}}} d(x, \partial\Omega)^p d\mu_n = \mu_n^{[p]}(\Omega \setminus \overline{K_{M,r}})$ is uniformly small: $(\mu_n^{[p]})_n$ is tight. As we have

$$\begin{aligned} \mu_n^{[p]}(\Omega) &= \text{Pers}_p(\mu_n) = \text{OT}_p^p(\mu_n, 0) \\ &\leq (\text{OT}_p(\mu_n, \nu) + \text{OT}_p(\nu, 0))^p \rightarrow (\text{OT}_p(\mu, \nu) + \text{OT}_p(\nu, 0))^p, \end{aligned}$$

it is also bounded in total variation. Hence, $(\mu_n^{[p]})_n$ is relatively compact for the weak convergence: this concludes the proof. \square

Proof of Lemma 4.9. Let $P = \sum_{i=1}^N \lambda_i \delta_{a_i}$ be a probability distribution with $a_i \in \mathcal{D}_f$ of mass $m_i \in \mathbb{N}$, and define $m_{\text{tot}} = \sum_{i=1}^N m_i$. By Proposition 4.4, every Fréchet mean a of P is in correspondence with a Fréchet mean for the Wasserstein distance \tilde{a} of $\tilde{P} = \sum_{i=1}^N \lambda_i \delta_{\tilde{a}_i}$, where $\tilde{a}_i = a_i + (m_{\text{tot}} - m_i) \delta_{\partial\Omega}$, with a being the restriction of \tilde{a} to Ω .

Fix $m \in \mathbb{N}$, and let $\tilde{a}_1, \dots, \tilde{a}_N$ be point measures of mass m in $\tilde{\Omega}$. Write $\tilde{a}_i = \sum_{j=1}^m \delta_{x_{i,j}}$, so that $x_{i,j} \in \tilde{\Omega}$ for $1 \leq i \leq N$, $1 \leq j \leq m$, with the $x_{i,j}$ non-necessarily distinct. Define

$$T : (x_1, \dots, x_N) \in \tilde{\Omega}^N \mapsto \arg \min \left\{ \sum_{i=1}^N \lambda_i \rho(x_i, y)^p, y \in \tilde{\Omega} \right\} \in \tilde{\Omega}. \quad (4.1)$$

Since we assume $p > 1$, T is well-defined and is continuous (the minimizer is unique by strict convexity). Using the localization property stated in [COO15a, Section 2.2], we know that the support of a Fréchet mean of \tilde{P} is included in the finite set

$$S := \{T(x_{1,j_1}, \dots, x_{N,j_N}), 1 \leq j_1, \dots, j_N \leq m\}.$$

Let $K = m^N$ and let z_1, \dots, z_K be an enumeration of the points of S (with potential repetitions). Denote by $\text{Gr}(z_k)$ the N elements x_1, \dots, x_N , with $x_i \in \text{spt}(\tilde{a}_i)$, such that $z_k = T(x_1, \dots, x_N)$. It is explained in [COO15a, §2.3], that finding a Fréchet mean of \tilde{P} is equivalent to finding a minimizer of the problem

$$\inf_{(\gamma_1, \dots, \gamma_N) \in \Pi} \sum_{i=1}^N \lambda_i \iint_{\tilde{\Omega}^2} \rho(x_i, y)^p d\gamma_i(x_i, y), \quad (4.2)$$

where Π is the set of plans $(\gamma_i)_{i=1, \dots, N}$, with γ_i having for first marginal \tilde{a}_i , and such that all γ_i s share the same (non-fixed) second marginal. Furthermore, we can assume without loss of generality that $(\gamma_1 \dots \gamma_N)$ is supported on $(\text{Gr}(z_k), z_k)_k$, i.e. a point z_k in the Fréchet mean is necessary transported to its corresponding grouping $\text{Gr}(z_k)$ by (optimal) $\gamma_1, \dots, \gamma_N$ [COO15a, §2.3]. For such a minimizer, the common second marginal is a Fréchet mean of \tilde{P} .

A potential minimizer of (4.2) is described by a vector $\gamma = (\gamma_{i,j,k}) \in \mathbb{R}_+^{NmK}$ such that:

$$\begin{cases} \text{for } 1 \leq i \leq N, 1 \leq j \leq m, & \sum_{k=1}^K \gamma_{i,j,k} = 1 \text{ and} \\ \text{for } 2 \leq i \leq N, 1 \leq k \leq K, & \sum_{j=1}^m \gamma_{i,j,k} = \sum_{j=1}^m \gamma_{1,j,k}. \end{cases} \quad (4.3)$$

Let $c \in \mathbb{R}^{NmK}$ be the vector defined by $c_{i,j,k} = \mathbf{1}\{x_{i,j} \in \text{Gr}(z_k)\} \lambda_i \rho(x_{i,j}, z_k)^p$. Then, the problem (4.2) is equivalent to

$$\underset{\gamma \in \mathbb{R}_+^{NmK}}{\text{minimize}} \gamma^T c \quad \text{under the constraints (4.3)}. \quad (4.4)$$

The set of Fréchet means of P are in bijection with the set of minimizers of this Linear Programming problem (see [Sch03, §5.15]), which is given by a face of the polyhedron described by the equations (4.3). Hence, if we show that this polyhedron is integer (i.e. its vertices have integer values), then it would imply that the extreme points of the set of Fréchet means of P are point measures, concluding the proof. The constraints (4.3) are described by a matrix A of size $(Nm + (N-1)K) \times NmK$ and a vector $b = [\mathbf{1}_{Nm}, \mathbf{0}_{(N-1)K}]$, such that $\gamma \in \mathbb{R}^{NmK}$ satisfies (4.3) if and only if $A\gamma = b$.

A sufficient condition for the polyhedron $\{Ax \leq b\}$ to be integer is to satisfy the following property (see [Sch03, Section 5.17]): for all $u \in \mathbb{Z}^{NmK}$, the dual problem

$$\max\{y^T b, y \geq 0 \text{ and } y^T A = u\} \quad (4.5)$$

has either no solution (i.e. there is no $y \geq 0$ satisfying $y^T A = u$), or it has an integer optimal solution y .

For y satisfying $y^T A = u$, write $y = [y^0, y^1]$ with $y^0 \in \mathbb{R}^{Nm}$ and $y^1 \in \mathbb{R}^{(N-1)K}$, so that y^0 is indexed on $1 \leq i \leq N$, $1 \leq j \leq m$ and y^1 is indexed on $2 \leq i \leq N$, $1 \leq k \leq K$. One can check that, for $2 \leq i \leq N$, $1 \leq j \leq m$, $1 \leq k \leq K$:

$$u_{1,j,k} = y_{1,j}^0 + \sum_{i'=2}^N y_{i',k}^1 \quad \text{and} \quad u_{i,j,k} = y_{i,j}^0 - y_{i,k}^1, \quad (4.6)$$

so that,

$$\begin{aligned} y^T b &= \sum_{i=1}^N \sum_{j=1}^m y_{i,j}^0 = \sum_{j=1}^m y_{1,j}^0 + \sum_{i=2}^N \sum_{j=1}^m y_{i,j}^0 \\ &= \sum_{j=1}^m (u_{1,j,k} - \sum_{i=2}^N y_{i,k}^1) + \sum_{i=2}^N \sum_{j=1}^m (u_{i,j,k} + y_{i,k}^1) \\ &= \sum_{i=1}^N \sum_{j=1}^m u_{i,j,k}. \end{aligned}$$

Therefore, the function $y^T b$ is constant on the set $P := \{y \geq 0, y^T A = u\}$, and any point of the set is an argmax. We need to check that if the set P is non-empty, then it contains a vector with integer coordinates: this would conclude the proof. A solution of the homogeneous equation $y^T A = 0$ satisfies $y_{i,j}^0 = y_{i,k}^1 = \lambda_i$ for $i \geq 2$ and $y_{1,j}^0 = -\sum_{i=2}^N y_{i,k}^1 = -\sum_{i=2}^N \lambda_i$ and reciprocally, any choice of $\lambda_i \in \mathbb{R}$ gives rise to a solution of the homogeneous equation. For a given u , one can verify that the set of solutions of $y^T A = u$ is given, for $\lambda_i \in \mathbb{R}$, by

$$\begin{cases} y_{1,j}^0 = \sum_{i=1}^N u_{i,j,k} - \sum_{i=2}^N \lambda_i \\ y_{i,j}^0 = \lambda_i \text{ for } i \geq 2, \\ y_{i,k}^1 = -u_{i,j,k} + \lambda_i \text{ for } i \geq 2. \end{cases}$$

Such a solution exists if and only if for all j , $U_j := \sum_{i=1}^N u_{i,j,k}$ does not depend on k and for $i \geq 2$, $U_{i,k} := u_{i,j,k}$ does not depend on j . For such a

vector u , P corresponds to the $\lambda_i \geq 0$ with $\lambda_i \geq \max_k U_{i,k}$ and $U_j \geq \sum_{i=1}^N \lambda_i$. If this set is non empty, it contains as least the point corresponding to $\lambda_i = \max\{0, \max_k U_{i,k}\}$, which is an integer: this point is integer valued, concluding the proof. \square

Part II

**Applications in statistics and
learning**

CHAPTER 5

FAST ESTIMATION OF FRÉCHET MEANS FOR PERSISTENCE DIAGRAMS

Abstract

In this chapter, we address the problem of estimating Fréchet means (a.k.a. barycenters) in the space of persistence diagrams from a computational perspective. A theoretical study of Fréchet means of persistence diagrams and measures can be found in Chapter 4, where existence results are provided. Here, we propose an approach to approximate Fréchet means of persistence diagrams in a convex, parallelizable, and GPU-friendly way. This work is essentially taken from [LCO18], although being significantly more detailed and revisited under the light of Chapter 3.

Implementation resources.

- [Implementation in Gudhi of Wasserstein distances and Lagrangian barycenters for persistence diagrams.](#)
- (To appear in [Gudhi](#)) Eulerian barycenters using entropic regularization.
- [A tutorial in Gudhi for barycenters of persistence diagrams.](#)

5.1 Preliminary remarks and problem formulation

Let Ω denote the open half-plane

$$\Omega := \{(b, d) \in \mathbb{R}^2, b > d\}, \quad (5.1)$$

and let $\partial\Omega$ denote its boundary, namely the diagonal

$$\partial\Omega := \{(t, t), t \in \mathbb{R}\}. \quad (5.2)$$

Eventually, let $\bar{\Omega}$ denote $\Omega \sqcup \partial\Omega$.

For now, we fix a parameter $p \in [1, +\infty)$. Recall from Definition 2.9 in Section 2.1 that persistence diagrams are defined as locally finite point measures supported on Ω , that is of the form $\sum_{x \in X} n_x \delta_x$, where $X \subset \Omega$ is locally finite, and $n_x \in \mathbb{N}$ for all $x \in X$. Recall that persistence diagrams can be equipped with metrics \mathbf{d}_p , defining a metric space $(\mathcal{D}^p, \mathbf{d}_p)$, and which depends on a ground metric d which is set to be $d(x, y) := \|x - y\|_p$ in this chapter. The space of persistence diagrams is denoted by \mathcal{D} . It is a subset of the space of non-negative Radon measures supported on Ω , denoted by $\mathcal{M}(\Omega)$ or \mathcal{M} if there is no ambiguity.

Here, as we target numerical applications, we will only consider finite sets of finite persistence diagrams. Therefore, we can assume that our diagrams are supported on $[0, 1]^2 \cap \Omega$.

Problem formulation In this context, we consider the following problem. Given a set of (finite) persistence diagrams $\nu_1, \dots, \nu_N \in \mathcal{D}^p$, and non-negative weights $(\lambda_i)_{1 \leq i \leq N}$ that sum to one, we are looking for minimizers of

$$\mathcal{E} : \mu \mapsto \sum_{i=1}^N \lambda_i \mathbf{d}_p(\mu, \nu_i)^p. \quad (5.3)$$

This problem was initially introduced in [TMMH14], we refer the reader to Section 2.1.3 for more details. This problem is known to have (non-unique, in general) solutions in both \mathcal{D} and \mathcal{M} (and the minimum values are the same). However, these exact solutions are generally intractable. The goal is thus to provide a way to estimate such minimizers in practice, ideally in a large-scale manner.

Diagram metrics as an optimal transport problems (reminder). In this context, following the results of Chapter 3 (namely, Proposition 3.5

and Proposition 3.15), the metric \mathbf{d}_p between two persistence diagrams $\mu = \sum_{i=1}^{n_1} \delta_{x_i}$ and $\nu = \sum_{j=1}^{n_2} \delta_{y_j}$ can be reformulated as a *balanced* optimal transport problem (see Section 2.2) on the space $\tilde{\Omega}$ obtained from $\bar{\Omega}$ by identifying all the points in $\partial\Omega$ to a single point (still denoted by $\partial\Omega$). First, note that d naturally induces a ground cost on $\tilde{\Omega} \times \tilde{\Omega}$ (still denoted by d) by setting $d(x, \{\partial\Omega\}) := d(x, \mathbf{p}_{\partial\Omega}(x))$, where $\mathbf{p}_{\partial\Omega}(x)$ is the orthogonal projection of $x \in \Omega$ onto the diagonal $\partial\Omega$.

Introduce the operator

$$\begin{aligned} \mathbf{R} : \mathcal{M}(\Omega) &\rightarrow \mathcal{M}(\{\partial\Omega\}) \\ \mu &\mapsto \mu(\Omega)\delta_{\partial\Omega}, \end{aligned} \quad (5.4)$$

and let $\tilde{\mu} = \mu + \mathbf{R}\nu$ and $\tilde{\nu} = \nu + \mathbf{R}\mu$. Let $a = (\mathbf{1}_{n_1}, n_2) \in \mathbb{R}^{n_1+1}$ and $b = (\mathbf{1}_{n_2}, n_1) \in \mathbb{R}^{n_2+1}$ be the weight vectors of μ and ν respectively, where $\mathbf{1}_k = (1 \dots 1) \in \mathbb{R}^k$. Let C be the cost matrix with block structure

$$C = \begin{pmatrix} \hat{C} & u \\ v^T & 0 \end{pmatrix} \in \mathbb{R}^{(n_1+1) \times (n_2+1)}, \quad (5.5)$$

with $u_i = d(x_i, \partial\Omega)^p$, $v_j = d(\partial\Omega, y_j)^p$ and $\hat{C} = (d(x_i, y_j)^p)_{ij}$.

It reads

$$\mathbf{d}_p(\mu, \nu)^p = \mathbf{L}_C(a, b) \quad (5.6)$$

where \mathbf{L}_C is defined by the following optimization problems, dual of each other:

$$\mathbf{L}_C(a, b) = \min_{P \in \Pi(a, b)} \langle C, P \rangle = \min_{P \in \Pi(a, b)} \sum_{ij} C_{ij} P_{ij}, \quad (5.7)$$

$$= \max_{f, g \in \Psi_C} \langle a, f \rangle + \langle b, g \rangle = \max_{f, g \in \Psi_C} \sum_i a_i f_i + \sum_j b_j g_j, \quad (5.8)$$

where the constraint sets are respectively defined as

$$\Pi(a, b) := \{P \in \mathbb{R}^{(n_1+1) \times (n_2+1)}, P\mathbf{1}_{n_2+1} = a, P^T\mathbf{1}_{n_1+1} = b\}, \quad (5.9)$$

$$\Psi_C := \{(f, g) \in \mathbb{R}^{n_1+1} \times \mathbb{R}^{n_2+1}, \forall i, j, f_i + g_j \leq C_{ij}\}. \quad (5.10)$$

REMARK 5.1. Note that the cost d defined on $\tilde{\Omega} \times \tilde{\Omega}$ is not a metric as one could have $d(x, y) \geq d(x, \partial\Omega) + d(\partial\Omega, y)$. To recover a true metric space, we can define

$$\rho(x, y) := \min\{d(x, y), d(x, \partial\Omega) + d(\partial\Omega, y)\}.$$

It is proved (Proposition 3.15) that using either ρ or d leads to the same optimal transport plans and optimal transport costs, so actually these two costs are interchangeable. Formally, it yields with previous notations,

$$\mathbf{L}_C(a, b) = W_{p, \rho}^p(\hat{\mu}, \hat{\nu}).$$

From a computational perspective, using $d(x, y)^p = \|x - y\|_p^p$, which is a separable cost, is highly beneficial when implementing the entropic regularization of diagram distances.

5.2 A Lagrangian approach

The first algorithm proposed to address the problem (5.3), assuming $p = 2$, was also introduced in [TMMH14]. It relies on a Lagrangian approach and is summarized in Algorithm 1. As it performs barycenter estimation relying on the Munkres (Hungarian) algorithm, it is referred to as the *B-Munkres algorithm* in the following. This algorithm is proved to converge to a *local* minimum of \mathcal{E} in \mathcal{D} , see Figure 5.1 for an illustrative example. In this section, we propose to revisit this algorithm by adopting an optimal transport perspective.

For $x_1 \dots x_m \in \tilde{\Omega}$, assume that $x_1 \dots x_k \in \Omega$ (i.e. are off-diagonal points) and $x_{k+1} \dots x_m = \partial\Omega$, let \bar{x} denote the arithmetic mean of $x_1 \dots x_k$, and define

$$\text{mean}(x_1 \dots x_m) := \arg \min_x \sum_{i=1}^n d(x, x_i)^2 = \frac{k}{m} \bar{x} + \frac{m-k}{m} \mathbf{p}_{\partial\Omega}(\bar{x}).$$

Interestingly, under the light of (5.6), Algorithm 1 can formally be retrieved as a slight variation of algorithms introduced independently in the optimal transport literature at a similar time in [CD14]. To clarify this, we introduce the following proposition.

PROPOSITION 5.2. *Let ν_1, \dots, ν_N be N persistence diagrams of respective masses (number of points) n_1, \dots, n_N . Let $m_{\text{tot}} := \sum_i n_i$. Define $w_{n_i} := (\mathbf{1}_{n_i}, m_{\text{tot}} - n_i) \in \mathbb{R}^{n_i+1}$. Minimizing the energy*

$$\mathcal{E} : \mu \mapsto \sum_{i=1}^N \mathbf{d}_p^p(\mu, \nu_i)$$

is equivalent to minimizing

$$\begin{aligned} \tilde{\mathcal{E}} : (\tilde{\Omega})^{m_{\text{tot}}} &\rightarrow \mathbb{R} \\ (y_1 \dots y_{m_{\text{tot}}}) &\mapsto \mathbf{L}_{C_i}(w_k, w_{n_i}) \end{aligned}$$

Algorithm 1 Estimation of Fréchet means (B-Munkres) [TMMH14]

Input: Set of diagrams $\nu_1 \dots \nu_N$, initial guess $Y = \{y_1 \dots y_k\}$.
Output: A persistence diagram $\mu = \sum_j \delta_{y_j}$ that is a local minima of \mathcal{E} on \mathcal{D} with $p = 2$.
 Let converged = False.
while not converged do
 $Z \leftarrow \emptyset$ /* store new points arising from $\partial\Omega$ */.
 for $i = 1 \dots N$ **do**
 Compute γ_i optimal partial matching between μ and ν_i ,
 $Z \leftarrow Z \cup \left\{ \frac{x}{m} + \frac{(m-1)}{m} \mathbf{p}_{\partial\Omega}(x), (\mathbf{p}_{\partial\Omega}(x), x) \in \gamma_i \right\}$.
 end for
 for each $y_j \in \mu$ **do**
 Let $x_j^i \in \text{spt}(\nu_i) \cup \partial\Omega$ be such that $(y_j, x_j^i) \in \gamma_i$
 $y_j \leftarrow \text{mean}(x_j^i, 1 \leq i \leq N)$
 end for
 if None of the $(y_j)_j$ has changed and Z is empty **then**
 converged \leftarrow True.
 else
 $Y \leftarrow Y \cup Z$.
 end if
end while
return μ .

where $k := |\{i, y_i \in \Omega\}|$ (number of off-diagonal points) and C_i is the pairwise distance matrix defined in (5.5) (between the off-diagonal points of $\mu := \sum_{i, y_i \in \Omega} \delta_{y_i}$ and ν_i). Furthermore, if $(y_1 \dots y_{m_{\text{tot}}})$ is a minimizer of $\tilde{\mathcal{E}}$, then μ is a minimizer of \mathcal{E} , and vice-versa.

Proof. In Chapter 4, we proved (Proposition 4.4 and Lemma 4.9) that minimizing \mathcal{E} on \mathcal{D} was equivalent to minimize

$$\mathcal{F} : \mathcal{D}_{m_{\text{tot}}}(\tilde{\Omega}) \rightarrow \mathbb{R}$$

$$\tilde{\mu} \mapsto \sum_{i=1}^N W_{p,\rho}^p(\tilde{\mu}, \nu_i + (m_{\text{tot}} - n_i)\delta_{\partial\Omega}),$$

where $\mathcal{D}_{m_{\text{tot}}}(\tilde{\Omega})$ denotes the measures of the form $\tilde{\mu} = \sum_{j=1}^{m_{\text{tot}}} \delta_{y_j}$, where $y_j \in \tilde{\Omega}$, in sense that if $\tilde{\mu}$ minimizes \mathcal{F} , then $\mu := \sum_{i, y_i \in \Omega} \delta_{y_j}$ minimizes \mathcal{E} , and vice-




Figure 5.1: Example of output of Algorithm 1.

versa.

Furthermore, in this context (see Remark 5.1),

$$W_{p,\rho}^p(\tilde{\mu}, \nu_i + (m_{\text{tot}} - n_i)\delta_{\partial\Omega}) = \mathbf{L}_{C_i}(w_k, w_{n_i}),$$

proving the claim. \square

This problem is *exactly* a standard optimal-transport barycenter problem. A first remark that can be done from this crucial observation is that we immediately obtain that minimizers of $\tilde{\mathcal{E}}$ (thus of \mathcal{E} once we restrict to off-diagonal points) must have their support included in

$$\{\text{mean}(x_1 \dots x_N), x_1 \in (a_1 \cup \{\partial\Omega\}) \dots x_N \in (a_N \cup \{\partial\Omega\})\},$$

a property known as the localization property [COO15a, Section 2.2].

In [CD14, §4.3], Cuturi and Doucet propose to consider the problem of minimizing $\tilde{\mathcal{E}}$ as a minimization problem over the positions $\{y_1 \dots y_{m_{\text{tot}}}\}$ describing (the current estimation of) the barycenter. They derive a Newton update rule that reads (in the context of persistence diagrams)

$$\{y_1 \dots y_{m_{\text{tot}}}\} = Y \leftarrow \text{mean}(X_i P_i^{*,T}),$$

where X_i denotes the positions of the points of a_i (completed with a mass $m_{\text{tot}} - a_i(\Omega)$ on $\partial\Omega$), and $P_i^{*,T}$ is an optimal transport plan between Y and X_i for the cost matrix $C(X, Y)$ defined in (5.5). Enforcing P_i to be a permutation matrix (which is always possible thanks to the Birkhoff-von Neumann theorem, see Section 2.2) exactly yields the update rule of Algorithm 1. To sum up, Algorithm 1 can be understood as a particular instance¹ of [CD14, Algorithm 2].

Limitations. In particular, Algorithm 1 only provides a local minimizer of \mathcal{E} in \mathcal{D} . Here, a local minimizer must be understood as an element $\mu \in \mathcal{D}$ for which there is some $\varepsilon > 0$, such that

$$\forall \nu \in \mathcal{D}, d_2(\mu, \nu) \leq \varepsilon \Rightarrow \mathcal{E}(\mu) \leq \mathcal{E}(\nu).$$

One can naturally wonder if such local minimizers do exist, and how “bad” they are. Unfortunately, we prove that \mathcal{E} admits arbitrary bad local minimizers in \mathcal{D} , to which the algorithm converges when initialized with one of the input diagrams (as suggested in [TMMH14]).

LEMMA 5.3. *For any $\kappa \geq 1$, there exists a set of diagrams ν_1, \dots, ν_N such that, when initialized with one of the $(\nu_i)_i$, the B-Munkres algorithm converges to a diagram μ_{loc} such that*

$$\mathcal{E}(\mu_{\text{loc}}) \geq \kappa \mathcal{E}^*$$

where $\mathcal{E}^* = \min_{\mu \in \mathcal{D}} \mathcal{E}(\mu)$.

Proof. To alleviate notations, we write in this proof the coordinates of the points in Ω with respect to the axes given by $\partial\Omega, \partial\Omega^\perp$. Fix an integer $N \in \mathbb{N}$ and let $\varepsilon_N = \frac{1}{2N}$. For $1 \leq i \leq N$, let ν_i be a diagram with N points $x_1^{(i)}, \dots, x_N^{(i)}$ of respective coordinates $((i-1)N + k)\varepsilon_N, 1$, see Figure 5.2 for an illustration with $N = 3$. Let also ν_{N+1} be the empty diagram.

Assume that Algorithm 1, when applied to estimate a Fréchet mean of $X_N := \nu_1 \dots \nu_{N+1}$ with weights $\lambda_1 = \dots = \lambda_{N+1} = \frac{1}{N+1}$, is initialized on ν_{N+1} , so that initially the estimate is described by $Y = \emptyset$. Let \mathcal{E}_N denote the corresponding energy function.

The optimal partial matching γ_{N+1} between the current estimate and ν_{N+1} is empty, while we have for $1 \leq i \leq N$,

$$\gamma_i = \{(\mathbf{p}_{\partial\Omega}(x_k^{(i)}), x_k^{(i)}), k = 1 \dots N\}.$$

¹More precisely, a slight variation, as Cuturi and Doucet only considered the Euclidean cost. On the quotient space $\tilde{\Omega}$, d is not the Euclidean cost, but computations adapt faithfully thanks to the block structure of C mentioned in (5.5).




Figure 5.2: Illustration of the diagrams used in the proof with $N = 3$. Purple crosses represent the output of the algorithm when initialized on the empty diagram, denoted by μ_{loc} in the proof. Orange squares represent the competitor $\hat{\mu}$ that performs, asymptotically, arbitrarily better than μ_{loc} .

Thus, the update step consists in adding to Y the N^2 points

$$y_j = \left((i-1)N + k, \frac{1}{N+1} \right), \quad \text{where } j = (i-1)N + k, 1 \leq j \leq N^2.$$

At the second iteration, the optimal partial matchings are respectively given by

$$\gamma_i = \{(y_j, x_k^{(i)}) \text{ if } j = (i-1)N + k, (y_j, \mathbf{p}_{\partial\Omega}(y_j)) \text{ otherwise}\}.$$

In particular, no new point needs to be added. Furthermore, one can check that the update rule let Y unchanged, reaching the stopping criterion of the algorithm. As an output of Algorithm 1, this resulting $\mu_{\text{loc}} = \sum_{j=1}^{N^2} \delta_{y_j}$ is necessarily a local minimum of \mathcal{E} on \mathcal{D} [TMMH14, §3.1]. The energy reached by this diagram is

$$\begin{aligned} \mathcal{E}_N(\mu_{\text{loc}}) &= \frac{1}{N+1} N^2 \cdot \underbrace{\left(\frac{1}{N+1} \right)^2}_{\mathbf{d}_2^2(\mu_{\text{loc}}, \nu_{N+1})} \\ &\quad + \frac{N}{N+1} \cdot \underbrace{\left(N \cdot \left(\frac{N}{N+1} \right)^2 + N(N-1) \left(\frac{1}{N+1} \right)^2 \right)}_{\mathbf{d}_2^2(\mu_{\text{loc}}, \nu_i)}, \end{aligned}$$

so that essentially $\mathcal{E}_N(\mu_{\text{loc}}) \sim N$ as $N \rightarrow \infty$.

On the other hand, let us introduce the diagram with N points $\hat{\mu} := \sum_{j=1}^N \delta_{\hat{y}_j}$, with

$$\hat{y}_j := \text{mean}((x_j^{(i)})_{i=1}^N, \partial\Omega) = \left(((j-1)N + (N-1)/2) \varepsilon_N, \frac{N}{N+1} \right).$$

The optimal partial matching between $\hat{\mu}$ and ν_i , for $1 \leq i \leq N$, is given by

$$\hat{\gamma}_i = \{(\hat{y}_j, x_j^{(i)}), j = 1 \dots N\},$$

inducing a cost of

$$\mathbf{d}_2^2(\hat{\mu}, \nu_i) = N \cdot \left(\left(\frac{N+1}{2} - i \right)^2 \varepsilon_N^2 + \left(\frac{1}{N+1} \right)^2 \right),$$

while the distance to $\nu_{N+1} = \mathbf{0}$ is simply $N \left(\frac{N}{N+1} \right)^2$. Eventually, the energy of $\hat{\mu}$ is

$$\mathcal{E}_N(\hat{\mu}) = \frac{N}{N+1} \left(\frac{N}{N+1} \right)^2 + \frac{N}{N+1} \sum_{i=1}^N \left(\left(\frac{N+1}{2} - i \right)^2 \varepsilon_N^2 + \left(\frac{1}{N+1} \right)^2 \right),$$

so that, since $\varepsilon_N = 2^{-N}$, $\mathcal{E}_N^* \leq \mathcal{E}_N(\hat{\mu}) \rightarrow 1$ as $N \rightarrow \infty$, and in particular $\mathcal{E}_N(\mu_{\text{loc}})/\mathcal{E}_N^*$ can be arbitrarily large. \square

Aside from its non-convexity, Algorithm 1 suffers in practice from a high computational cost. Indeed, solving partial matching problems between two diagrams of size n scales (just as any optimal transport problem) essentially in $\mathcal{O}(n^3)$. As Algorithm 1 requires, at each iteration, to compute N optimal partial matchings, it is untractable when the number of diagram N and the typical size of the diagram n increase, which is confirmed numerically (see Figure 5.8 below).

Improving on these two issues. Algorithm 1 produces an estimation of a Fréchet mean by optimizing on positions: it is a *Lagrangian* approach (see Section 2.2). Lagrangian approaches are known to lead to non-convex optimization problems (see for instance [CD14, §4.4]). In contrast, we adopt below a *Eulerian* approach which—at the cost of discretizing the ground space—allows us to recover a convex optimization problem.

In terms of computational efficiency, we rely on an *entropic regularization* of optimal transport, see Section 2.2.4 for an introduction. A careful adaptation of computational techniques developed in computational optimal transport [PC17] yields to a scalable problem than can be run efficiently on GPUs, drastically improving running times on large sets of large diagrams.

REMARK 5.4. *We end this section by mentioning a concurrent approach to improve on Algorithm 1 introduced in [VBT19, KVT19]. Briefly speaking, it consists in using the Auction algorithm—as adapted in [KMN17]—to estimate optimal partial matchings between persistence diagrams. This approach significantly reduces the running time in Algorithm 1 and has the benefits of providing an output that is an actual persistence diagram (while the approach we develop below returns a measure supported on the discretized groundspace). However, the resulting problem is still not convex.*

5.3 Entropic approximation for distances and means of persistence diagrams

Following the correspondence established in (5.6), entropic regularization of optimal transport metrics (we refer the reader to Section 2.2.4 for more details) can be used to approximate the diagram distance $\mathbf{d}_p(\cdot, \cdot)$. Given two persistence diagrams μ, ν with respective mass n_1, n_2 , let $n := n_1 + n_2$, $a = (\mathbf{1}_{n_1}, n_2) \in \mathbb{R}^{n_1+1}$, $b = (\mathbf{1}_{n_2}, n_1) \in \mathbb{R}^{n_2+1}$, and let C be the cost matrix as defined in (5.5). In this context, given a regularization parameter $\gamma > 0$, one can approximate the distance between μ and ν using one of the following optimization problems, dual of each other:

$$\mathbf{d}_p(\mu, \nu)^p \simeq \mathbf{d}_p^\gamma(\mu, \nu)^p := \inf_{P \in \Pi(a, b)} \langle P, C \rangle + \gamma h(P), \quad (5.11)$$

$$= \sup_{(f, g) \in \mathbb{R}^{n_1+1} \times \mathbb{R}^{n_2+1}} \langle a, f \rangle + \langle b, g \rangle \quad (5.12)$$

$$- \varepsilon \sum_{ij} \exp\left(\frac{f_i + g_j - C_{ij}}{\varepsilon}\right), \quad (5.13)$$

$$=: \mathbf{L}_C^\gamma(a, b) \quad (5.14)$$

where $\Pi(a, b)$ is defined in (5.9), $h(P) := \sum_{ij} P_{ij}(\log(P_{ij}) - 1)$ is the *negentropy*. For an optimal P^γ , we define the *Sinkhorn cost* as $\mathbf{S}_C^\gamma(a, b) := \langle P^\gamma, C \rangle$. Note that one has $\mathbf{S}_C^\gamma(a, b) \geq \mathbf{L}_C^\gamma(a, b)$.

The seminal strength of the minimization problem involved in (5.11) is that it is strictly convex and thus admits a unique solution P^γ . This solution can be approximated by iterating the *Sinkhorn map*

$$S : (u, v) \in \mathbb{R}^{n_1+1} \times \mathbb{R}^{n_2+1} \mapsto \left(\frac{a}{Kv}, \frac{b}{K^T u} \right), \quad (5.15)$$

where $K = \exp(-C/\varepsilon) \in \mathbb{R}^{(n_1+1) \times (n_2+1)}$ (where exponentiation must be understood term-wise). Indeed, if (u^γ, v^γ) is a fixed point of (5.15), then

it turns out that $u^\gamma K v^\gamma = P^\gamma$. Furthermore, for any initialization (u_0, v_0) , the sequence defined by

$$u_{t+1} = \frac{a}{K v_t} \quad (5.16)$$

$$v_{t+1} = \frac{b}{K^T u_{t+1}}, \quad (5.17)$$

converges to an optimal pair (u^γ, v^γ) . Note also the relation $(u^\gamma, v^\gamma) = (\exp(f^\gamma/\gamma), \exp(g^\gamma/\gamma))$, where (f^γ, g^γ) is optimal in (5.14).

5.3.1 Error control

The above definition gives us a way to approximate \mathbf{d}_p by \mathbf{d}_p^γ by simply iterating the Sinkhorn map S (note that in practice, additional care is required to obtain much faster computations). Of course, one might ask for some control on the error made by using \mathbf{d}_p^γ instead of \mathbf{d}_p . There are two sources of error:

- The regularization error induced by the parameter γ
- The fact that, in practice, we do not exactly reach a fixed point of S as this would be required to perform “infinitely many” iterations.

Therefore, in practice, we obtain a transport plan $P_t^\gamma = \text{diag}(u_t^\gamma) K \text{diag}(v_t^\gamma)$ where (u_t^γ, v_t^γ) is the output after t iterations of the Sinkhorn map (5.15). The error made by approximating $\mathbf{d}_p(\mu, \nu)^p$ by $\langle P_t^\gamma, C \rangle$ has been studied in [AWR17]. Adapted to our context, their bounds read

$$|\mathbf{d}_p(\mu, \nu)^p - \langle P_t^\gamma, C \rangle| \leq 2\gamma n \log(n) + \text{dist}(P_t^\gamma, \Pi(a, b)) \|C\|_\infty \quad (5.18)$$

where

$$\text{dist}(P, \Pi(a, b)) := \|P\mathbf{1} - a\|_1 + \|P^T\mathbf{1} - b\|_1,$$

(that is, error on marginals). It remains to understand at which rate $\text{dist}(P, \Pi(a, b)) \rightarrow 0$ as the number of iterations $t \rightarrow \infty$. In [DGK18], authors prove that iterating the Sinkhorn map (5.15) gives a plan P_t^γ satisfying $\text{dist}(P_t^\gamma, \Pi(a, b)) < \varepsilon$ in $\mathcal{O}\left(\frac{\|C\|_\infty^2}{\gamma\varepsilon} + \ln(n)\right)$ iterations. Given (5.18), a natural choice is thus to take $\gamma = \frac{\varepsilon}{n \ln(n)}$ for a desired precision ε , which lead to a total of $\mathcal{O}\left(\frac{n \ln(n) \|C\|_\infty^2}{\varepsilon^2}\right)$ iterations in the Sinkhorn loop. These results can be used to pre-tune parameters t and γ to control the approximation error due to smoothing. However, these are worst-case bounds, controlled by max-norms, and are often too pessimistic in practice. To overcome this

phenomenon, we introduce on-the-fly error control, using approximate solutions to the smoothed primal (5.11) and dual (5.14) optimal transport problems, which provide upper and lower bounds on the optimal transport cost.

Upper and Lower Bounds. The Sinkhorn algorithm, after at least one iteration ($t \geq 1$), produces feasible variables for the (non-regularized) dual problem (5.8):

$$(f_t^\gamma, g_t^\gamma) = (\gamma \log(u_t^\gamma), \gamma \log(v_t^\gamma)) \in \Psi_C,$$

providing a natural lower bound of the target quantity

$$\langle f_t^\gamma, a \rangle + \langle g_t^\gamma, b \rangle \leq \sup_{f, g \in \Psi_C} \langle a, f \rangle + \langle b, g \rangle = \mathbf{d}_p(\mu, \nu)^p.$$

However, it turns out that in practice, this quantity performs poorly as a lower bound of the true optimal transport cost (see Figure 5.3 and Section 5.3.4 below) in most of our experiments. To improve on this, we computed the so called C -transform $(f_t^\gamma)^c$ of f_t^γ [San15, §1.6], defined as:

$$\forall j, (f_t^\gamma)_j^c = \max_i \{C_{ij} - f_i\}, \quad j \leq n_2 + 1.$$

Applying a C^T -transform on $(f_t^\gamma)^c$, we recover two vectors $(f_t^\gamma)^{c\bar{c}} \in \mathbb{R}^{n_1+1}$, $(f_t^\gamma)^c \in \mathbb{R}^{n_2+1}$. For any feasible f, g , we have that [PC17, Prop 3.1]

$$\langle f, a \rangle + \langle g, b \rangle \leq \langle f^{c\bar{c}}, a \rangle + \langle f^c, b \rangle.$$

When C 's top-left block is the squared Euclidean metric, this problem can be cast as that of computing the Moreau envelope [Mor65] of f . In a Eulerian setting and when our measures are described as histograms supported on a finite regular grid (which we consider below), the C -transform in linear time with respect to the grid resolution d by using either a linear-time Legendre transform or the Parabolic Envelope algorithm [Luc10, §2.2.1, §2.2.2].

Unlike dual iterations, the *primal* iterate P_t^γ does *not* belong to the transport polytope $\Pi(a, b)$ after a finite number t of iterations. We use the `rounding_to_feasible` algorithm introduced in [AWR17] to compute efficiently a feasible approximation R_t^γ of P_t^γ that does belong to $\Pi(a, b)$. Putting these two elements together, we obtain

$$\underbrace{\langle (f_t^\gamma)^{c\bar{c}}, a \rangle + \langle (f_t^\gamma)^c, b \rangle}_{m_t^\gamma} \leq \mathbf{L}_C(a, b) \leq \underbrace{\langle R_t^\gamma, C \rangle}_{M_t^\gamma}. \quad (5.19)$$




Figure 5.3: (Left) $M_t^\gamma := \langle R_t^\gamma, C \rangle$ (red) and $m_t^\gamma := \langle \alpha_t^{c\bar{c}}, a \rangle + \langle \alpha_t^c, b \rangle$ (green) as a function of t , the number of iterations of the Sinkhorn map (t ranges from 1 to 500, with fixed $\gamma = 10^{-3}$). (Middle) Final M^γ (red) and m^γ (green) provided by Alg.1, computed for decreasing γ s, ranging from 10^{-1} to $5 \cdot 10^{-4}$. For each value of γ , Sinkhorn loop is run until $d(P_t^\gamma, \Pi(a, b)) < 10^{-3}$. Note that the γ -axis is flipped. (Right) Influence of $c\bar{c}$ -transform for the Sinkhorn dual cost. (orange) The dual cost $\langle \alpha_t^\gamma, a \rangle + \langle \beta_t^\gamma, b \rangle$, where $(\alpha_t^\gamma, \beta_t^\gamma)$ are Sinkhorn dual variables (before the C -transform). (green) Dual cost after C -transform, i.e. with $((\alpha_t^\gamma)^{c\bar{c}}, (\alpha_t^\gamma)^c)$. Experiment run with $\gamma = 10^{-3}$ and 500 iterations.

Therefore, after iterating the Sinkhorn map (5.15) t times, we have that if $M_t^\gamma - m_t^\gamma$ is below a certain criterion ε , then we can guarantee that $\langle R_t^\gamma, C \rangle$ is *a fortiori* an ε -approximation of $\mathbf{L}_C(a, b)$. Observe that one can also have a relative error control: if one has $M_t^\gamma - m_t^\gamma \leq \varepsilon M_t^\gamma$, then $(1 - \varepsilon)M_t^\gamma \leq L_C(a, b) \leq M_t^\gamma$. Note that m_t^γ might be negative but can always be replaced by $\max(m_t^\gamma, 0)$ since C has non-negative entries (and therefore $\mathbf{L}_C(a, b) \geq 0$), while M_t^γ is always non-negative.

5.3.2 Computational aspects

Discretization. As mentioned in the beginning of the section, we assume in the remaining that our diagrams have their support in $[0, 1]^2 \cap \Omega$. As explained previously, updating positions (Lagrangian approaches) generally leads to non-convex problem. To recover a convex problem, we propose instead to adopt an Eulerian approach: just as a persistence diagram μ can be describe as a set of positions $\{x_1 \dots x_n\}$, it can also be seen as a map $x \in \Omega \mapsto 1$ if x is a point of the diagram, 0 otherwise. That is, we will optimize the mass put on each x in Ω . Of course, this theoretically leads to an infinite-dimensional optimization problem. In practice, we have to discretize the ground space to recover a tractable optimization problem, that has the huge advantage of being convex.

From a numerical perspective, encoding persistence diagrams as histograms on the square offers numerous advantages. Given a uniform grid of size $d \times d$ on $[0, 1]^2$, we associate to a given diagram μ a *matrix-shaped*

histogram $\mathbf{a} \in \mathbb{R}^{d \times d}$ such that \mathbf{a}_{ij} is the number of points in μ belonging to the cell located at position (i, j) in the grid (we use bold-faced small letters to insist on the fact that these histograms must be stored as square matrices). Note that, in order to account for the mass located on $\partial\Omega$, we must add an additional coordinate that will be denoted by Δ . We extend the operator \mathbf{R} to histograms, associating to a histogram $\mathbf{a} \in \mathbb{R}^{d \times d}$ its total mass on the Δ coordinate (the $(d^2 + 1)$ -th one).

Convolutions. In the Eulerian setting, where diagrams are matrix-shaped histograms of size $d \times d = d^2$, the cost matrix would have size $d^2 \times d^2$. Since one may want to use large values of d to have low discretization error (typically $d = 100$), instantiating the cost matrix naively is usually intractable. However, [SDGP⁺15] showed that for regular grids endowed with a separable cost, each Sinkhorn iteration (as well as other key operations such as evaluating Sinkhorn's cost $\mathbf{S}_C^\gamma = \langle P^\gamma, C \rangle$) can be performed using Gaussian convolutions, which amounts to performing matrix multiplications of size $d \times d$, without having to manipulate $d^2 \times d^2$ matrices at any step. The framework we consider is slightly different due to the extra dimension $\{\partial\Omega\}$ (and the corresponding row and column in (5.5)), but equivalent computational properties hold. This observation is crucial from a numerical perspective. The idea is sketched in Figure 5.4.

Implementation details. Let (\mathbf{u}, u_Δ) be a pair where $\mathbf{u} \in \mathbb{R}^{d \times d}$ is a matrix-shaped histogram and $u_\Delta \in \mathbb{R}_+$ is a real number accounting for the mass located the Δ coordinate (i.e. on $\{\partial\Omega\}$). We denote by $\vec{\mathbf{u}}$ the $d^2 \times 1$ column vector obtained when reshaping \mathbf{u} . The $(d^2 + 1) \times (d^2 + 1)$ cost matrix C and corresponding kernel K are given by

$$C = \begin{pmatrix} \widehat{C} & \vec{\mathbf{c}}_\Delta \\ \vec{\mathbf{c}}_\Delta^T & 0 \end{pmatrix}, \quad K = \begin{pmatrix} \widehat{K} := e^{-\frac{\widehat{C}}{\gamma}} & \vec{\mathbf{k}}_\Delta := e^{-\frac{\vec{\mathbf{c}}_\Delta}{\gamma}} \\ \vec{\mathbf{k}}_\Delta^T & 1 \end{pmatrix},$$

where $\widehat{C} = (\|(i, i') - (j, j')\|_p^p)_{ii', jj'}$, $\mathbf{c}_\Delta = (\|(i, i') - \pi_\Delta((i, i'))\|_p^p)_{ii'}$. The crucial point is that C and K as defined above will never be instantiated, because we can rely instead on $\mathbf{c} \in \mathbb{R}^{d \times d}$ defined as $\mathbf{c}_{ij} = |i - j|^p$ and $\mathbf{k} = e^{-\frac{\mathbf{c}}{\gamma}}$.

PROPOSITION 5.5 (Iteration of Sinkhorn map). *The operation of K to (\mathbf{u}, u_Δ) can be performed as:*

$$(\mathbf{u}, u_\Delta) \mapsto (\mathbf{k}(\mathbf{u}\mathbf{k}^T)^T + u_\Delta \mathbf{k}_\Delta, \langle \mathbf{u}, \mathbf{k}_\Delta \rangle + u_\Delta) \quad (5.20)$$

where $\langle \cdot, \cdot \rangle$ denotes the Froebenius dot-product in $\mathbb{R}^{d \times d}$.

We now introduce $\mathbf{m} := \mathbf{k} \odot \mathbf{c}$ and $\mathbf{m}_\Delta := \mathbf{k}_\Delta \odot \mathbf{c}_\Delta$ (\odot denotes term-wise multiplication).

PROPOSITION 5.6 (Computation of \mathbf{S}_C^γ). *Let $(\mathbf{u}, u_\Delta), (\mathbf{v}, v_\Delta) \in \mathbb{R}^{d \times d+1}$. The transport cost of $P := \text{diag}(\vec{\mathbf{u}}, u_\Delta)K\text{diag}(\vec{\mathbf{v}}, v_\Delta)$ can be computed as:*

$$\langle \text{diag}(\vec{\mathbf{u}}, u_\Delta)K\text{diag}(\vec{\mathbf{v}}, v_\Delta), C \rangle = \langle \text{diag}(\vec{\mathbf{u}})\widehat{K}\text{diag}(\vec{\mathbf{v}}), \widehat{C} \rangle + u_\Delta \langle \mathbf{v}, \mathbf{m}_\Delta \rangle + v_\Delta \langle \mathbf{u}, \mathbf{m}_\Delta \rangle, \quad (5.21)$$

where the first term can be computed as:

$$\langle \text{diag}(\vec{\mathbf{u}})\widehat{K}\text{diag}(\vec{\mathbf{v}}), \widehat{C} \rangle = \|\mathbf{u} \odot (\mathbf{m}(\mathbf{kv}^T)^T + \mathbf{k}(\mathbf{mv}^T)^T)\|_1. \quad (5.22)$$

The last proposition addresses the computation of the upper bound involved in (5.19). First, we recall the `rounding_to_feasible` algorithm introduced by Altschuler et al. in [AWR17]. Here, $r(P)$ and $c(P)$ denote respectively the first and second marginal of a matrix P .

Algorithm 2 Rounding algorithm of [AWR17]

- 1: **Input:** $P \in \mathbb{R}^{d \times d}$, desired marginals r, c .
 - 2: **Output:** $F(P) \in \Pi(r, c)$ close to P .
 - 3: $X = \min\left(\frac{r}{r(P)}, 1\right) \in \mathbb{R}^d$
 - 4: $P' = \text{diag}(X)P$
 - 5: $Y = \min\left(\frac{c}{c(P')}, 1\right) \in \mathbb{R}^d$
 - 6: $P'' = P'\text{diag}(Y)$
 - 7: $e_r = r - r(P'')$, $e_c = c - c(P'')$
 - 8: **return** $F(P) := P'' + e_r e_c^T / \|e_c\|_1$
-

Finally, consider two histograms $(\mathbf{a}, a_\Delta), (\mathbf{b}, b_\Delta) \in \mathbb{R}^{d \times d} \times \mathbb{R}$, let $R \in \Pi((\mathbf{a}, a_\Delta), (\mathbf{b}, b_\Delta))$ be the rounded matrix of P as introduced in Algorithm 2. Let $r(P), c(P) \in \mathbb{R}^{d \times d} \times \mathbb{R}$ denote the first and second marginal of P respectively. We introduce (minimum and divisions must be understood term-wise):

$$X = \min\left(\frac{(\mathbf{a}, a_\Delta)}{r(P)}, \mathbf{1}\right), \quad Y = \min\left(\frac{(\mathbf{b}, b_\Delta)}{c(\text{diag}(X)P)}, \mathbf{1}\right),$$

along with $P' = \text{diag}(X)P\text{diag}(Y)$ and the marginal errors:

$$(e_r, (e_r)_\Delta) = (\mathbf{a}, a_\Delta) - r(P'), \quad (e_c, (e_c)_\Delta) = (\mathbf{b}, b_\Delta) - c(P'),$$




Figure 5.4: Sketch of the convolution technique used to efficiently iterate the Sinkhorn map (5.15) in a Eulerian setting.

PROPOSITION 5.7 (Computation of the upper bound $\langle R, C \rangle$). *The transport cost induced by R can be computed as:*

$$\begin{aligned} \langle R, C \rangle &= \langle \text{diag}(X \odot (\mathbf{u}, u_\Delta)) K \text{diag}(Y \odot (\mathbf{v}, v_\Delta)), C \rangle \\ &\quad + \frac{1}{\|e_c\|_1 + (e_c)_\Delta} (\|e_r^T \mathbf{c} e_c\|_1 + \|e_r \mathbf{c} e_c^T\|_1) \\ &\quad + (e_c)_\Delta \langle e_r, \mathbf{c}_\Delta \rangle + (e_r)_\Delta \langle e_c, \mathbf{c}_\Delta \rangle. \end{aligned} \quad (5.23)$$

Note that the first term can be computed using (5.21)

The proofs of these propositions are purely technical and can be found in Section 5.3.5.

Parallelization and GPU. Using a Eulerian representation is particularly beneficial when applying Sinkhorn’s algorithm, as shown by [Cut13]. Indeed, the Sinkhorn map (5.15) only involves matrix-vector operations. Furthermore, the cost matrix C is constant: it does not depend on the pair of histograms involved. When dealing with a large number of histograms, concatenating these histograms and running Sinkhorn’s iterations in parallel as matrix-matrix product results in significant speedup that can exploit GPGPU to compare a large number of pairs simultaneously. This makes this approach especially well-suited for large sets of persistence diagrams. This idea is sketched in Figure 5.5.

Algorithm 3 allows us to estimate distances between persistence diagrams in parallel by performing only $(d \times d)$ -sized matrix multiplications, leading to a computational scaling in d^3 where d is the grid resolution parameter. Note that a standard stopping threshold in Sinkhorn iteration process is to check the error to marginals $\text{dist}(P, \Pi(\mathbf{a}, \mathbf{b}))$, as motivated by (5.18).




Figure 5.5: Sketch of parallel computations of S when dealing with multiple pairs of diagrams.

Algorithm 3 Sinkhorn cost for persistence diagrams

Input: Pairs of PDs $(D_i, D'_i)_i$, smoothing parameter $\gamma > 0$, grid step $d \in \mathbb{N}$, stopping criterion, initial (\mathbf{u}, \mathbf{v}) .

Output: Approximation of all $(\mathbf{d}_p(D_i, D'_i)^p)_i$, upper and lower bounds if wanted.

init Cast D_i, D'_i as histograms $\mathbf{a}_i, \mathbf{b}_i$ on a $d \times d$ grid

while stopping criterion not reached **do**

Iterate in parallel (5.15) $(\mathbf{u}, \mathbf{v}) \mapsto S(\mathbf{u}, \mathbf{v})$ using (5.20)

end while

Compute all $\mathbf{S}_C^\gamma(\mathbf{a}_i + \mathbf{R}\mathbf{b}_i, \mathbf{b}_i + \mathbf{R}\mathbf{a}_i)$ using (5.21)

if Want a upper bound **then**

Compute $\langle R_i, C \rangle$ in parallel using (5.23)

end if

if Want a lower bound **then**

Compute $\langle (\alpha_i^\gamma)^{c\bar{c}}, \mathbf{a}_i \rangle + \langle (\alpha_i^\gamma)^c, \mathbf{b}_i \rangle$ using [Luc10]

end if

5.3.3 Entropic regularization for Fréchet means in \mathcal{D}

In addition to numerical efficiency, a major advantage of the regularized optimal transport is that

$$a \mapsto \mathbf{L}_C^\gamma(a, b)$$

is differentiable. In the Eulerian setting, where C is fixed, its gradient is given by centering the vector $\gamma \log(u^\gamma)$ where u^γ is a fixed point of the Sinkhorn map (5.15) (see [CD14]). This result can be adapted to our framework, namely:

PROPOSITION 5.8. *Let $\nu_1 \dots \nu_N$ be PDs, and $(\mathbf{a}_i)_i$ the corresponding histograms on a $d \times d$ grid. Let C_d be the corresponding $(d^2 + 1) \times (d^2 + 1)$ cost*




Figure 5.6: Barycenter estimation for different γ s with a simple set of 3 PDs (red, blue and green). The smaller the γ , the better the estimation (\mathcal{E} decreases, note the γ -axis is flipped on the right plot), at the cost of more iterations in Alg. 4. The mass appearing along the diagonal is a consequence of entropic smoothing: it does not cost much to delete while it increases the entropy of transport plans.

matrix. Note that C_d does not depend on \mathbf{x} . The gradient of the functional

$$\mathcal{E}^{\gamma,d} : \mathbf{x} \mapsto \sum_{i=1}^N \mathbf{L}_{C_d}^{\gamma}(\mathbf{x} + \mathbf{R}\mathbf{a}_i, \mathbf{a}_i + \mathbf{R}\mathbf{x})$$

is given by

$$\nabla_{\mathbf{x}} \mathcal{E}^{\gamma,d} = \gamma \left(\sum_{i=1}^N \log(u_i^{\gamma}) + \mathbf{R}^T \log(v_i^{\gamma}) \right) \quad (5.24)$$

where \mathbf{R}^T denotes the adjoint operator \mathbf{R} and $(u_i^{\gamma}, v_i^{\gamma})$ is a fixed point of the Sinkhorn map obtained while transporting $\mathbf{x} + \mathbf{R}\mathbf{a}_i$ onto $\mathbf{a}_i + \mathbf{R}\mathbf{x}$ for the cost C .

This result follows from the envelope theorem, taking into account the fact that \mathbf{x} appears in both terms depending on u and v . This formula can be exploited to compute barycenters via gradient descent, yielding Algorithm 4. Following [CD14, §4.2], we use a multiplicative update. This is a particular case of mirror descent [BT03] and is equivalent to a (Bregman) projected gradient descent on the positive orthant, retaining positive coefficients (masses) throughout iterations.

REMARK 5.9 (A comparison with linear representations). *When doing statistical analysis with persistence diagrams, a standard approach is to transform a diagram into a finite-dimensional vector—in a stable way—and then perform statistical analysis and learning with a Euclidean structure. This approach does not preserve the Wasserstein-like geometry of the diagram space and thus loses the algebraic interpretability of persistence diagrams. Figure 5.7 gives a qualitative illustration of the difference between Wasserstein barycenters (Fréchet mean) of persistence diagrams and*

Algorithm 4 Estimation of Fréchet means of persistence diagrams

Input: Persistence diagrams ν_1, \dots, ν_N , learning rate λ , smoothing parameter $\gamma > 0$, grid step $d \in \mathbb{N}$.

Output: Estimated barycenter \mathbf{x}

Init: \mathbf{x} uniform measure above the diagonal.

Cast each ν_i as a histogram \mathbf{a}_i on a $d \times d$ grid

while \mathbf{x} changes **do**

 Iterate S defined in (5.15) in parallel between all the pairs $(\mathbf{x} + \mathbf{R}\mathbf{a}_i)_i$ and $(\mathbf{a}_i + \mathbf{R}\mathbf{x})_i$, using (5.20).

$\nabla := \gamma(\sum_i \log(u_i^\gamma) + \mathbf{R}^T \log(v_i^\gamma))$

$\mathbf{x} := \mathbf{x} \odot \exp(-\lambda \nabla)$

end while

if Want energy **then**

 Compute $\frac{1}{N} \sum_i \mathbf{S}_C^\gamma(\mathbf{x} + \mathbf{R}\mathbf{a}_i, \mathbf{a}_i + \mathbf{R}\mathbf{x})$ using (5.21)

end if

Return \mathbf{x}

Euclidean barycenters (linear means) of persistence images [AEK⁺17], a commonly used vectorization for PDs [MKPY16, ZZJS16, OHK18].

REMARK 5.10. *Aside from its computational strengths (convexity, parallelism, GPU-capable), this Eulerian approach has another major advantage over Lagrangian ones: it straightforwardly extends to any measure supported on $[0, 1]^2 \cap \Omega$ and is not restricted to persistence diagram (nor even to measures with finite support). Therefore, this approach would allow us to compute (more precisely, to estimate) distance and barycenters between more complicated measures supported on Ω , such as expected persistence diagrams (see Section 7.1).*

5.3.4 Experiments

Let us now showcase Algorithm 3 and Algorithm 4 in numerical applications. In the following, experiments are run with $p = 2$, but would work with any finite $p \geq 1$. The choice of $p = 2$ is consistent with the work of [TMMH14] for Fréchet mean estimation.

Fast barycenters Let us compare Algorithm 4 (referred to as *Sinkhorn*) with the combinatorial algorithm of the B-Munkres algorithm Algorithm 1. Figure 5.8 records running times of both algorithms on a set of 10 diagrams having from n to $2n$ points, n ranging from 1 to 500, on Intel Xeon 2.3




Figure 5.7: Illustration of differences between Fréchet means with Wasserstein and Euclidean geometry. The top row represents input data, namely persistence diagrams (left), vectorization of diagrams as persistence images in $\mathbb{R}^{100 \times 100}$ (middle, [AEK⁺17]), and discretization of diagrams as histograms (right). The bottom row represents the estimated barycenters (orange scale) with input data (shaded), using the approach of Algorithm 1 [TMMH14] (left), the arithmetic mean of persistence images (middle) and our optimal transport based approach (right).

GHz (CPU) and P100 (GPU, Sinkhorn only). When running Algorithm 4, the gradient descent is performed until $|\mathcal{E}(\mathbf{x}_{t+1})/\mathcal{E}(\mathbf{x}_t) - 1| < 0.01$, with $\gamma = 10^{-1}/n$ and $d = 50$. Our experiment shows that Algorithm 4 drastically outperforms *B-Munkres* as the number of points n increases. We interrupt *B-Munkres* at $n = 30$, after which computational time becomes an issue.

Aside the computational efficiency, we highlight the benefits of operating with a convex formulation in Figure 5.9. Due to non-convexity, the *B-Munkres* algorithm is only guaranteed to converge to a local minima, and its output depends on initialization. We illustrate on a toy set of $N = 3$ diagrams how our algorithm avoids local minima thanks to the Eulerian approach we take.

***k*-means on large PD sets.** We now merge Algorithm 3 and Algorithm 4 in order to perform unsupervised clustering via *k*-means on PDs. We work with the 3D-shape database provided by [SP04] and generate diagrams in the same way as in [COO15b], working in practice with 5000 diagrams




Figure 5.8: Average running times for B-Munkres (*blue*) and Sinkhorn (*red*) algorithms (log-log scale) to average 10 PDs.




Figure 5.9: Qualitative comparison of B-Munkres and Algorithm 4. (a) Input set of $N = 3$ diagrams with $n = 20$ points each. (b) Output of B-Munkres when initialized on the blue diagram (orange squares) and input data (grey scale). (c) Output of B-Munkres initialized on the green diagram. (d) Output of Algorithm 4 on a 100×100 grid, $\gamma = 5 \cdot 10^{-4}$, learning-rate $\lambda = 5$, Sinkhorn stopping criterion (distance to marginals): 0.001, gradient descent performed until $|\mathcal{E}(\mathbf{z}_{t+1})/\mathcal{E}(\mathbf{z}_t) - 1| < 0.01$.—As one can see, localization of masses is similar. Initialization of B-Munkres is made on one of the input diagram as indicated in [TMMH14, Alg. 1], and leads to convergence to different local minima. Our convex approach (Algorithm 4) performs better (lower energy). As a baseline, the energy of the naive arithmetic mean of the three diagrams is 0.72.




Figure 5.10: Illustration of our k -means algorithm. From left to right: 20 diagrams extracted from *horses* and *camels* plot together (one color for each diagram); the centroid they are matched with provided by our algorithm; 20 diagrams of *head* and *faces*; along with their centroid; decrease of the objective function. Running time depends on many parameters along with the random initialization of k -means. As an order of magnitude, it takes from 40 to 80 minutes with this 5000 PD dataset on a P100 GPU.

with 50 to 100 points each. The database contains 6 classes: *camel*, *cat*, *elephant*, *horse*, *head* and *face*. In practice, this unsupervised clustering algorithm detects two main clusters: faces and heads on one hand, camels and horses on the other hand are systematically grouped together. Figure 5.10 illustrates the convergence of our algorithm and the computed centroids for the aforementioned clusters.

5.3.5 Proofs of Propositions 5.5, 5.6, 5.7

Proof of Proposition 5.5. Given a histogram $\mathbf{u} \in \mathbb{R}^{d \times d}$ and a mass $u_\Delta \in \mathbb{R}_+$, one can observe that:

$$\widehat{K}\mathbf{u} = \mathbf{k}(\mathbf{k}\mathbf{u}^T)^T. \quad (5.25)$$

In particular, the operation $\mathbf{u} \mapsto \widehat{K}\mathbf{u}$ can be performed by only manipulating matrices in $\mathbb{R}^{d \times d}$.

Indeed, observe that:

$$\widehat{K}_{ij,kl} = e^{-(i-k)^2/\gamma} e^{-(j-l)^2/\gamma} = \mathbf{k}_{ik}\mathbf{k}_{jl},$$

so we have:

$$\begin{aligned} (\widehat{K}\mathbf{u})_{i,j} &= \sum_{k,l} K_{ij,kl} \mathbf{u}_{k,l} \\ &= \sum_{k,l} \mathbf{k}_{ik}\mathbf{k}_{jl} \mathbf{u}_{k,l} = \sum_k \mathbf{k}_{ik} \sum_l \mathbf{k}_{jl} \mathbf{u}_{kl} \\ &= \sum_k \mathbf{k}_{ik} (\mathbf{k}\mathbf{u}^T)_{jk} = (\mathbf{k}(\mathbf{k}\mathbf{u}^T)^T)_{i,j}. \end{aligned}$$

Thus we have in our case:

$$K(\mathbf{u}, u_\Delta) = (\widehat{K}\mathbf{u} + u_\Delta \mathbf{k}_\Delta, \langle \mathbf{u}, \mathbf{k}_\Delta \rangle + u_\Delta)$$

where $\langle a, b \rangle$ designs the Froebenius dot product between two histograms $a, b \in \mathbb{R}^{d \times d}$. These computations only involves manipulation of matrices with size $d \times d$. \square

Proof of Proposition 5.6.

$$\begin{aligned} \langle \text{diag}(\vec{\mathbf{u}}) \widehat{K} \text{diag}(\vec{\mathbf{v}}), \widehat{C} \rangle &= \sum_{ijkl} \mathbf{u}_{ij} \mathbf{k}_{ik} \mathbf{k}_{jl} [\mathbf{c}_{ik} + \mathbf{c}_{jl}] \mathbf{v}_{kl} \\ &= \sum_{ijkl} \mathbf{u}_{ij} ([\mathbf{k}_{ik} \mathbf{c}_{ik}] \mathbf{k}_{jl} \mathbf{v}_{kl} + \mathbf{k}_{ik} [\mathbf{k}_{jl} \mathbf{c}_{jl}] \mathbf{v}_{kl}) \\ &= \sum_{ij} \mathbf{u}_{ij} \sum_{kl} (\mathbf{m}_{ik} \mathbf{k}_{jl} \mathbf{v}_{kl} + \mathbf{k}_{ik} \mathbf{m}_{jl} \mathbf{v}_{kl}) \end{aligned}$$

Thus, we finally have:

$$\langle \text{diag}(\vec{\mathbf{u}}) \widehat{K} \text{diag}(\vec{\mathbf{v}}), \widehat{C} \rangle = \|\mathbf{u} \odot (\mathbf{m}(\mathbf{k}\mathbf{v}^T)^T + \mathbf{k}\mathbf{m}\mathbf{v}^T)^T\|_1$$

And finally, taking the $\{\Delta\}$ bin into considerations,

$$\begin{aligned} &\langle \text{diag}(\vec{\mathbf{u}}, u_\Delta) K \text{diag}(\vec{\mathbf{v}}, v_\Delta), C \rangle \\ &= \left\langle \begin{pmatrix} \text{diag}(\vec{\mathbf{u}}) \widehat{K} \text{diag}(\vec{\mathbf{v}}) & v_\Delta (\vec{\mathbf{u}} \odot \vec{\mathbf{k}}_\Delta) \\ u_\Delta (\vec{\mathbf{v}}^T \odot \vec{\mathbf{k}}_\Delta^T) & u_\Delta v_\Delta \end{pmatrix}, \begin{pmatrix} \widehat{C} & \vec{\mathbf{c}}_\Delta \\ \vec{\mathbf{c}}_\Delta^T & 0 \end{pmatrix} \right\rangle \\ &= \langle \text{diag}(\vec{\mathbf{u}}) \widehat{K} \text{diag}(\vec{\mathbf{v}}), \widehat{C} \rangle + u_\Delta \langle \mathbf{v}, \mathbf{k}_\Delta \odot \mathbf{c}_\Delta \rangle + v_\Delta \langle \mathbf{u}, \mathbf{k}_\Delta \odot \mathbf{c}_\Delta \rangle, \end{aligned}$$

concluding the proof. \square

Proof of Proposition 5.7. By straightforward computations, the first and second marginals of $P_t^\gamma = \text{diag}(\vec{\mathbf{u}}) K \text{diag}(\vec{\mathbf{v}})$ are given by:

$$\left(\sum_{kl} \mathbf{u}_{ij} K_{ij,kl} \mathbf{v}_{kl} \right)_{ij} = \mathbf{u} \odot (K\mathbf{v}), \quad \left(\sum_{ij} \mathbf{u}_{ij} K_{ij,kl} \mathbf{v}_{kl} \right)_{kl} = (\mathbf{u}K) \odot \mathbf{v}.$$

Observe that $K\mathbf{v}$ and $\mathbf{u}K$ can be computed using Proposition 5.5.

Now, the transport cost computation is:

$$\begin{aligned} \langle F(P_t^\gamma), C \rangle &= \langle \text{diag}(X) P_t^\gamma \text{diag}(Y), C \rangle + \langle e_r e_c^T / \|e_c\|_1, C \rangle \\ &= \langle \text{diag}(X \odot \mathbf{u}) K \text{diag}(Y \odot \mathbf{v}), C \rangle + \frac{1}{\|e_c\|_1} \sum_{ijkl} (e_r)_{ij} (e_c)_{kl} [\mathbf{c}_{ik} + \mathbf{c}_{jl}] \end{aligned}$$

The first term is the transport cost induced by a rescaling of \mathbf{u}, \mathbf{v} and can be computed with Proposition 5.6. Consider now the second term. Without considering the additional bin $\{\Delta\}$, we have:

$$\begin{aligned} \sum_{ijkl} (e_r)_{ij} (e_c)_{kl} [\mathbf{c}_{ik} + \mathbf{c}_{jl}] &= \sum_{ijl} (e_r)_{ij} \sum_k \mathbf{c}_{ik} (e_c)_{kl} + \sum_{ijk} (e_r)_{ij} \sum_l \mathbf{c}_{jl} (e_c)_{kl} \\ &= \sum_{ijl} (e_r)_{ij} (\mathbf{c} e_c)_{il} + \sum_{ijk} (e_r)_{ij} (\mathbf{c} e_c^T)_{jk} \\ &= \|e_r^T \mathbf{c} e_c\|_1 + \|e_r \mathbf{c} e_c^T\|_1, \end{aligned}$$

so when we consider our framework (with $\{\Delta\}$), it comes:

$$\begin{aligned} \left\langle \begin{pmatrix} e_r \\ (e_r)_\Delta \end{pmatrix} \cdot (e_c \quad (e_c)_\Delta), C \right\rangle &= \left\langle \begin{pmatrix} e_r e_c^T & (e_c)_\Delta e_r \\ (e_r)_\Delta e_c^T & (e_r)_\Delta (e_c)_\Delta \end{pmatrix}, \begin{pmatrix} \widehat{C} & \vec{\mathbf{c}}_\Delta \\ \vec{\mathbf{c}}_\Delta^T & 0 \end{pmatrix} \right\rangle \\ &= \langle e_r e_c^T, \widehat{C} \rangle + (e_c)_\Delta \langle e_r, \mathbf{c}_\Delta \rangle + (e_r)_\Delta \langle e_c, \mathbf{c}_\Delta \rangle. \end{aligned}$$

Putting things together finally proves the claim. \square

CHAPTER 6

LINEAR REPRESENTATIONS OF PERSISTENCE DIAGRAMS

Abstract

In this chapter, we study linear representations of persistence diagrams, a common vectorization method to make use of persistence diagrams and measures in machine learning. The first section uses the characterization of convergence proved in Chapter 3 to derive a useful characterization of continuous linear representations in a very general theoretical framework. The second section proposes a neural network architecture that we call PERSLAY, which aims at *learning* a task-driven representation method in a supervised learning context. It is essentially based on [CCI⁺20].

Implementation resources.

- (To appear in [Gudhi](#)) implementation of PERSLAY. Preliminary version available [on Mathieu Carriere's github page](#), along with a [tutorial](#).

6.1 Continuity of linear representations

Recall that \mathcal{M}^p denotes the space of persistence measures (Radon measures supported on the open half-plane Ω with finite total persistence).




Figure 6.1: Some common linear representations of persistence diagrams. From left to right: A persistence diagram. Its persistence surface [AEK⁺17], which is a persistence measure. The corresponding persistence silhouette [CFL⁺14]. The corresponding Betti Curve [Ume17].

A *linear representation* of persistence measures (in particular persistence diagrams) is a map

$$\begin{aligned} \Phi : \mathcal{M}^p &\rightarrow \mathcal{B} \\ \mu &\mapsto \mu(f) \end{aligned}$$

for some Banach space \mathcal{B} and some chosen function $f : \Omega \rightarrow \mathcal{B}$. Doing so, one can turn a sample of diagrams (or measures) into a sample of vectors, making the use of machine learning tools easier. Of course, a natural requirement is that Φ must be continuous. In practice, building a linear representation (see below for a list of examples) generally follows the same pattern: first consider a “nice” function ϕ , e.g. a Gaussian distribution, then introduce a weight with respect to the distance to the diagonal $d(\cdot, \partial\Omega)^p$, and prove that $\mu \mapsto \mu(\phi(\cdot)d(\cdot, \partial\Omega)^p)$ has some regularity properties (continuity, stability, etc.). By relying on a result proved in Chapter 3, namely Theorem 3.7, we show that this approach always gives a continuous linear representation, and that it is the only way to do so.

6.1.1 Characterization of continuity

For $(\mathcal{B}, \|\cdot\|_{\mathcal{B}})$ a Banach space (typically \mathbb{R}^d), define the class of functions:

$$\mathcal{C}_{b,p}^0 = \left\{ f : \Omega \rightarrow \mathcal{B}, f \text{ continuous and } x \mapsto \frac{\|f(x)\|_{\mathcal{B}}}{d(x, \partial\Omega)^p} \text{ bounded} \right\} \quad (6.1)$$

PROPOSITION 6.1. *Let \mathcal{B} be a Banach space and $f : \Omega \rightarrow \mathcal{B}$ a function. The linear representation $\Phi : \mathcal{M}^p \rightarrow \mathcal{B}$ defined by $\Phi : \mu \mapsto \mu(f) = \int_{\Omega} f(x)d\mu(x)$ is continuous with respect to OT_p and $\|\cdot\|_{\mathcal{B}}$ if and only if $f \in \mathcal{C}_{b,p}^0$.*

Before proving this result, we recall the content of Theorem 3.7 introduced in Chapter 3: let $(\mu_n)_n$ be a sequence in \mathcal{M}^p , and $\mu \in \mathcal{M}^p$. One has

$$\text{OT}_p(\mu_n, \mu) \rightarrow 0 \Leftrightarrow \begin{cases} \mu_n \xrightarrow{v} \mu, \\ \text{Pers}_p(\mu_n) \rightarrow \text{Pers}_p(\mu). \end{cases}$$

It follows that (Corollary 3.10) that $\text{OT}_p(\mu_n, \mu) \rightarrow 0 \Leftrightarrow \mu_n^{[p]} \xrightarrow{w} \mu^{[p]}$, where the measure $\mu^{[p]}$ is defined, for $A \subset \Omega$ (Borel), by

$$\mu^{[p]}(A) := \int_A d(x, \partial\Omega)^p d\mu(x).$$

This strong characterization of the convergence in the metric space $(\mathcal{M}^p, \text{OT}_p)$ gives us a powerful tool to immediately study continuity questions, as showcased in the following proof.

Proof. Let $f \in \mathcal{C}_{b,p}^0$ and $\mu, \mu_1, \mu_2 \dots \in \mathcal{M}^p$ be such that $\text{OT}_p(\mu_n, \mu) \rightarrow 0$. Following [Nie11, Theorem 2], it is sufficient to treat the case where f takes values in \mathbb{R} ($\mathcal{B} = \mathbb{R}$). Using Corollary 3.10, it means that $\mu_n^{[p]} \xrightarrow{w} \mu^{[p]}$, and thus that

$$\int_{\Omega} \frac{f(x)}{d(x, \partial\Omega)^p} d\mu_n^{[p]}(x) \rightarrow \int_{\Omega} \frac{f(x)}{d(x, \partial\Omega)^p} d\mu^{[p]}(x),$$

that is

$$\Phi(\mu_n) = \int_{\Omega} f(x) d\mu_n(x) \rightarrow \int_{\Omega} f(x) d\mu(x) = \Phi(\mu),$$

i.e. Φ is continuous with respect to OT_p .

Conversely, let $f : \Omega \rightarrow \mathcal{B}$. Assume first that f is not continuous at some $x \in \Omega$. There exists a sequence $(x_n)_n \in \Omega^{\mathbb{N}}$ such that $x_n \rightarrow x$ but $f(x_n) \not\rightarrow f(x)$. Let $\mu_n = \delta_{x_n}$ and $\mu = \delta_x$. We have $\text{OT}_p(\mu_n, \mu) \rightarrow 0$, but $\mu_n(f) = f(x_n) \not\rightarrow f(x_0) = \mu(f)$, so that the linear representation $\mu \mapsto \mu(f)$ cannot be continuous.

Then, assume that f is continuous but that $x \mapsto \frac{f(x)}{d(x, \partial\Omega)^p}$ is not bounded. Let thus $(x_n)_n \in \Omega^{\mathbb{N}}$ be a sequence such that $\left\| \frac{f(x_n)}{d(x_n, \partial\Omega)^p} \right\| \rightarrow +\infty$. Define the measure $\mu_n := \frac{1}{\|f(x_n)\|} \delta_{x_n}$. Observe that $\text{OT}_p(\mu_n, 0) = \frac{d(x_n, \partial\Omega)^p}{\|f(x_n)\|} \rightarrow 0$ by hypothesis. However, $\|\mu_n(f)\| = 1$ for all n , allowing to conclude once again that $\mu \mapsto \mu(f)$ cannot be continuous. \square

EXAMPLE 6.2. *Let us give some examples of such linear representations (which are thus continuous) commonly used in applications of TDA. Note that the following definitions do not rely on the fact that the input must be*

a persistence diagram and actually make sense for any persistence measure in \mathcal{M}^p . See Figure 6.1 for an illustration, in which computations are done with $p = 1$ (and $p' = 1$ for the weighted Betti curve).

- Persistence surface and its variations. Let $K : \mathbb{R}^2 \rightarrow \mathbb{R}$ be a nonnegative Lipschitz continuous bounded function (e.g. $K(x, y) = \exp\left(-\frac{\|x-y\|^2}{2}\right)$) and define $f : x \in \Omega \mapsto d(x, \partial\Omega)^p \times K(x, \cdot)$, so that $f(x) : \mathbb{R}^2 \rightarrow \mathbb{R}$ is a real-valued function. The corresponding representation Φ takes its values in $(C_b(\mathbb{R}^2), \|\cdot\|_\infty)$, the (Banach) space of continuous bounded functions. This representation is called the persistence surface and has been introduced with slight variations in different works [AEK⁺17, CWRW15, KHF16, RHBK15].
- Persistence silhouettes. Let $\Lambda(x, t) = \max\left(\frac{x_2-x_1}{2} - \left|t - \frac{x_1+x_2}{2}\right|, 0\right)$ for $t \in \mathbb{R}$ and $x \in \Omega$. Then, defining $f : x \in \Omega \mapsto d(x, \partial\Omega)^{p-1} \times \Lambda(x, \cdot)$, one has that $\|f(x)\|_\infty$ is proportional to $d(x, \partial\Omega)^p$, so that the corresponding representation is continuous for OT_p . This representation is called the persistence silhouette, and was introduced in [CFL⁺14]. In particular, it consists in a weighted sum of the different functions of the persistence landscape [Bub15, BD17]. The corresponding Banach space is $(C_b(\mathbb{R}), \|\cdot\|_\infty)$.
- Weighted Betti curves. For $t \in \mathbb{R}$, define B_t the rectangle $(-\infty, t] \times [t, +\infty)$. Let $p, p' \geq 1$, and define $f : x \in \Omega \mapsto (t \mapsto d(x, \partial\Omega)^{p-1/p'} \mathbf{1}\{x \in B_t\})$. Then $f(x) \in L_{p'}(\mathbb{R})$ with $\|f(x)\|_{p'}$ proportional to $d(x, \partial\Omega)^p$. The corresponding function Φ is the weighted Betti curve, which takes its values in the Banach space $(L_{p'}(\mathbb{R}), \|\cdot\|_{p'})$. In particular, one obtains the continuity of the classical Betti curves from $(\mathcal{M}^1, \text{OT}_1)$ to $L_1(\mathbb{R})$.

6.1.2 Stability in the case $p = 1$.

Continuity is a basic expectation when embedding a set of diagrams (or measures) in some Banach space \mathcal{B} . One could however ask for more, e.g. some Lipschitz regularity: given a representation $\Phi : \mathcal{M}^p \rightarrow \mathcal{B}$, one may want to have $\|\Phi(\mu) - \Phi(\nu)\| \leq C \cdot \text{OT}_p(\mu, \nu)$ for some constant C . This property is generally referred to as “stability” in the TDA community and is generally obtained with $p = 1$, see for example [AEK⁺17, Theorem 5], [CCO17, Theorem 3.3 & 3.4], [STNR⁺18, §4], [RHBK15, Theorem 2], etc.

Here, we still consider the case of linear representations, and show that stability always holds with respect to the distance OT_1 . Informally, this is

explained by the fact that when $p = 1$, the cost function $(x, y) \mapsto d(x, y)^p$ is an actual distance.

PROPOSITION 6.3. *Define \mathcal{L} the set of Lipschitz continuous functions $f : \overline{\Omega} \rightarrow \mathbb{R}$ with Lipschitz constant less than 1 and that satisfy $f(\partial\Omega) = 0$. Let $T \subset \mathbb{R}$, and consider a family $(f_t)_{t \in T}$ with $f_t \in \mathcal{L}$. Then the linear representation $\Phi : \mu \mapsto (\mu(f_t))_{t \in T}$ is 1-Lipschitz continuous in the following sense:*

$$\|\Phi(\mu) - \Phi(\nu)\|_\infty := \sup_{t \in T} |(\mu - \nu)(f_t)| \leq \text{OT}_1(\mu, \nu), \quad (6.2)$$

for any measures $\mu, \nu \in \mathcal{M}^1$.

Proof. Consider $\mu, \nu \in \mathcal{M}^1$, and $\pi \in \text{Opt}(\mu, \nu)$ an optimal transport plan. Let $t \in T$. We have:

$$\begin{aligned} (\mu - \nu)(f_t) &= \int_{\Omega} f_t(x) d\mu(x) - \int_{\Omega} f_t(y) d\nu(y) = \iint_{\overline{\Omega} \times \overline{\Omega}} (f_t(x) - f_t(y)) d\pi(x, y) \\ &\leq \iint_{\overline{\Omega} \times \overline{\Omega}} d(x, y) d\pi(x, y) = \text{OT}_1(\mu, \nu), \end{aligned}$$

and thus, $\|\Phi(\mu) - \Phi(\nu)\|_\infty \leq \text{OT}_1(\mu, \nu)$. \square

REMARK 6.4. *As observed in Section 3.4, when $p = 1$, one actually has the Kantorovich-Rubinstein duality result $\text{OT}_1(\mu, \nu) = \sup_f |\mu(f) - \nu(f)|$, where the supremum is taken over 1-Lipschitz functions f that satisfy $f(\partial\Omega) = 0$.*

6.2 Learning representations of persistence diagrams using PersLay

In this section, as we target numerical applications, we restrict ourselves to finite persistence diagrams and will only consider linear representation valued in \mathbb{R}^d .

Linear representations are pretty handy: they are pretty simple to define and to compute. For instance, if $\mu = \sum n_x \delta_x$ is a persistence diagram, and $f : \Omega \rightarrow \mathbb{R}^d$, the representation $\mu(f)$ is simply given by $\sum n_x f(x) \in \mathbb{R}^d$. Therefore, given a set of persistence diagrams $\mu_1 \dots \mu_N$, one can recover a set $\mu_1(f) \dots \mu_N(f) \in \mathbb{R}^d$, on which basically every standard machine learning pipeline can be applied to solve a given task in a *supervised setting*. However, it might happen that the chosen function f is not adapted, or at least not optimal, to perform a given task; and there is no heuristic to choose f *a priori*.

6.2.1 Model

In [CCI⁺20], we considered the problem of optimizing the representation function f . Motivated by Section 6.1, we considered linear representations based on a function f of the form

$$x \mapsto w(x)\phi(x),$$

where $w : \mathbb{R}^2 \rightarrow \mathbb{R}$ is the *weight function* and $\phi : \mathbb{R}^2 \rightarrow \mathbb{R}^q$ is the *representation function*. The crucial point is that we consider *parameterized* functions $w = w_{\theta_1}$ and $\phi = \phi_{\theta_2}$, where $(\theta_1, \theta_2) \in \Theta \subset \mathbb{R}^D$ are *trainable* parameters that are learned during a training phase through a gradient descent.

REMARK 6.5. *There is some slight redundancy in using two functions w and ϕ . Formally, this helps to clarify how our formalism encompasses most vectorizations methods used in topological data analysis literature (also, it is implemented this way). In practice, it also has the benefits in terms of interpretability: the map $w : \mathbb{R}^2 \rightarrow \mathbb{R}$ can be easily visualized, giving a qualitative information on which areas of Ω are useful to achieve a good score in a given task (see Figure 6.4).*

REMARK 6.6. *Proposition 6.3 ensures that if $x \mapsto w(x)\phi(x)/d(x, \partial\Omega)^p$ is bounded, the vectorization $\mu \mapsto \mu(w\phi)$ would be continuous with respect to the OT_p metric. In practice, this would enforce to have to small values of w on points close to the diagonal. However, in supervised learning applications, there is no obvious reason for this, so we actually allow for more general (trainable) weight functions w .*

This flexible formalism allows us to recover and generalize standard vectorization methods used in topological data analysis. To that end we define our generic neural network layer for persistence diagrams, that we call PERSLAY. For a (finite) persistence diagram $\mu = \sum_{i=1}^n \delta_{x_i}$, it is defined through the following equation:

$$\text{PERSLAY}(\mu) := \mu(f) = \text{op}(\{w(x_i) \cdot \phi(x_i)\}_{1 \leq i \leq n}), \quad (6.3)$$

where op is any *permutation invariant operator* (such as minimum, maximum, sum, k th largest value...), $w : \mathbb{R}^2 \rightarrow \mathbb{R}$ is a *weight function* for the persistence diagram points, and $\phi : \mathbb{R}^2 \rightarrow \mathbb{R}^q$ is a *representation function* that we call *point transformation*, mapping each point (α_b, α_d) of a persistence diagram to a vector.

REMARK 6.7 (Relation to other works). *Note that this problem is also addressed in the concurrent work [HKN19], where authors consider a similar framework. They enforce $w = d(\cdot, \partial\Omega)^p$, and introduce essentially three point transformation functions ϕ , which roughly correspond to Gaussian, spike, and cone functions centered on the diagram points, and only consider $\text{op} = \text{sum}$. To that respect, our approach generalizes from their and allows our model to encompass a much larger class of vectorization methods used in practice. Furthermore, we focus on providing a publicly available implementation as a module of the `Gudhi` library, also hosted at <https://github.com/MathieuCarriere/perslay>.*

It is also important to note that both [CCI⁺20] and [HKN19] took some of their inspiration from [ZKR⁺17], introducing the notion of Deep Sets. In this work, authors proved that given a (countable, or compact) set X , a map $\Phi : X^n \rightarrow \mathbb{R}$ is permutation invariant—that is, it satisfies $\Phi(x_1 \dots x_n) = \Phi(x_{\sigma(1)} \dots x_{\sigma(n)})$ for any $(x_1 \dots x_n) \in \mathbb{X}^n$ and any permutation $\sigma \in \mathfrak{S}_n$ —if and only if it exists two maps $\rho_1 : X \rightarrow \mathbb{R}$ and $\rho_2 : \mathbb{R} \rightarrow \mathbb{R}$ such that $\Phi(x_1 \dots x_n) = \rho_2(\sum_{i=1}^n \rho_1(x_i))$. This expression is consistent with the one proposed by [HKN19] and with (6.3) with $\text{op} = \rho_2 \circ \sum$ and $w(\cdot)\phi(\cdot) = \rho_1(\cdot)$.

As explained above, in practice w and ϕ are of the form $w_{\theta_1}, \phi_{\theta_2}$ where the gradients of $\theta_1 \mapsto w_{\theta_1}$ and $\theta_2 \mapsto \phi_{\theta_2}$ are known and implemented so that back-propagation can be performed, and the parameters θ_1, θ_2 can be optimized during the training process. We emphasize that any neural network architecture ρ can be composed with PERSLAY to generate a neural network architecture for persistence diagrams. Let us now introduce three point transformation functions that we use and implement for parameter ϕ in (6.3).

- The *triangle point transformation*

$$\phi_{\Lambda} : \mathbb{R}^2 \rightarrow \mathbb{R}^q, x \mapsto [\Lambda_x(t_1), \Lambda_x(t_2), \dots, \Lambda_x(t_q)]^T,$$

where the triangle function Λ_x associated to a point $x = (b, d) \in \mathbb{R}^2$ is $\Lambda_x : t \mapsto \max\{0, d - |t - b|\}$, with $q \in \mathbb{N}$ and $t_1, \dots, t_q \in \mathbb{R}$.

- The *Gaussian point transformation*

$$\phi_{\Gamma} : \mathbb{R}^2 \rightarrow \mathbb{R}^q, x \mapsto [\Gamma_x(t_1), \Gamma_x(t_2), \dots, \Gamma_x(t_q)]^T,$$

where the Gaussian function Γ_x associated to a point $x = (b, d) \in \mathbb{R}^2$ is $\Gamma_x : t \mapsto \exp(-\|x - t\|_2^2 / (2\sigma^2))$ for a given $\sigma > 0$, $q \in \mathbb{N}$ and $t_1, \dots, t_q \in \mathbb{R}^2$.

- The *line point transformation*

$$\phi_L: \mathbb{R}^2 \rightarrow \mathbb{R}^q, x \mapsto [L_{\Delta_1}(x), L_{\Delta_2}(x), \dots, L_{\Delta_q}(x)]^T,$$

where the line function L_Δ associated to a line Δ with direction vector $e_\Delta \in \mathbb{R}^2$ and bias $b_\Delta \in \mathbb{R}$ is $L_\Delta: x \mapsto \langle x, e_\Delta \rangle + b_\Delta$, with $q \in \mathbb{N}$ and $\Delta_1, \dots, \Delta_q$ are q lines in the plane.

Formulation (6.3) is very general: despite its simplicity, it allows us to remarkably encode most persistence diagram vectorizations with a small set of point transformation functions ϕ , allowing to consider the choice of ϕ as a hyperparameter of sort. Let us show how it connects to most of the popular vectorizations and kernel methods for persistence diagrams in the literature.

- Using: $\phi = \phi_\Lambda$ with samples $t_1, \dots, t_q \in \mathbb{R}$, $\text{op} = k$ th largest value, $w = 1$ (a constant weight function), amounts to evaluating the *kth persistence landscape* [Bub15] on $t_1, \dots, t_q \in \mathbb{R}$.
- Using $\phi = \phi_\Lambda$ with samples $t_1, \dots, t_q \in \mathbb{R}$, $\text{op} = \text{sum}$, arbitrary weight function w , amounts to evaluating the *persistence silhouette* weighted by w [CFL⁺14] on $t_1, \dots, t_q \in \mathbb{R}$.
- Using $\phi = \phi_\Gamma$ with samples $t_1, \dots, t_q \in \mathbb{R}^2$, $\text{op} = \text{sum}$, arbitrary weight function w , amounts to evaluating the *persistence surface* weighted by w [AEK⁺17] on $t_1, \dots, t_q \in \mathbb{R}^2$. Moreover, characterizing points of persistence diagrams with Gaussian functions is also the approach advocated in several kernel methods for persistence diagrams [RHBK15, LY18, KHF16].
- Using $\phi = \phi_{\tilde{\Gamma}}$ where $\tilde{\Gamma}$ is a modification of the Gaussian point transformation defined with: $\tilde{\Gamma}_p = \Gamma_{\tilde{p}}$ for any $p = (x, y) \in \mathbb{R}^2$, where $\tilde{p} = p$ if $y \leq \nu$ for some $\nu > 0$, and $(x, \nu + \log(\frac{y}{\nu}))$ otherwise, $\text{op} = \text{sum}$, weight function $w = 1$, is the approach presented in [HKNU17].
- Using $\phi = \phi_L$ with lines $\Delta_1, \dots, \Delta_q \in \mathbb{R}^2$, $\text{op} = k$ th largest value, weight function $w = 1$, is similar to the approach advocated in [CCO17], where the sorted projections of the points onto the lines are then compared with the $\|\cdot\|_1$ norm and exponentiated to build the so-called Sliced Wasserstein kernel for persistence diagrams.




Figure 6.2: Some examples of orbits generated by the different choices of r (three simulations are represented for the different values of r).

6.2.2 Numerical experiments

We describe here some numerical experiments that were run in [CCI+20, §3.2, §4], and refer the reader to Appendix C therein for complementary details (e.g. report of the hyperparameters used). The accuracies reported in Table 6.1 and Table 6.2 represent a snapshot of the performances of some state-of-the-art methods at the time [CCI+20] was submitted.

Orbit classification. Our first application is on a synthetic dataset used as a benchmark in topological data analysis [LY18, CCO17, AEK+17]. It consists of sequences of points generated by different dynamical systems (see [HSW07]). Given some initial position $(x_0, y_0) \in [0, 1]^2$ and a parameter $r > 0$, we generate a point cloud $(x_n, y_n)_{n=1, \dots, N}$ following:

$$\begin{cases} x_{n+1} = x_n + ry_n(1 - y_n) & \text{mod } 1, \\ y_{n+1} = y_n + rx_{n+1}(1 - x_{n+1}) & \text{mod } 1. \end{cases} \quad (6.4)$$

The orbits of this dynamical system heavily depend on parameter r (see Figure 6.2). More precisely, for some values of r , voids might form in these orbits, and as such, persistence diagrams are likely to perform well at attempting to classify orbits with respect to the value of r generating them. As in previous works [LY18, CCO17, AEK+17], we use the five different parameters $r = 2.5, 3.5, 4.0, 4.1$ and 4.3 to simulate the different classes of orbits,

Dataset	PSS-K	PWG-K	SW-K	PF-K	PERSLAY
ORBIT5K	72.38(± 2.4)	76.63(± 0.7)	83.6(± 0.9)	85.9(± 0.8)	87.7(± 1.0)
ORBIT100K	—	—	—	—	89.2(± 0.3)

Table 6.1: Performance table. PSS-K, PWG-K, SW-K, PF-K stand for *Persistence Scale Space Kernel* [RHBK15], *Persistence Weighted Gaussian Kernel* [KHF16], *Sliced Wasserstein Kernel* [CCO17] and *Persistence Fisher Kernel* [LY18] respectively. We report the scores given in [LY18] for competitors on ORBIT5K, and the one we obtained using PERSLAY for both the ORBIT5K and ORBIT100K datasets. The latter dataset, and the corresponding diagrams, are too large in practice to be handled by kernel methods on persistence diagrams.

with random initialization of (x_0, y_0) and $N = 1,000$ points in each simulated orbit. These point clouds are then turned into persistence diagrams using a standard geometric filtration [CDSO14], called the *AlphaComplex* filtration¹ in dimensions 0 and 1. We generate two datasets: The first is ORBIT5K, where for each value of r , we generate 1000 orbits, ending up with a dataset of 5,000 point clouds. This dataset is the same as the one used in [LY18]. The second is ORBIT100K, which contains 20,000 orbits per class, resulting in a dataset of 100,000 point clouds—a scale that kernel for persistence diagrams cannot handle. This dataset aims to show the edge of our neural-network based approach over kernels methods when dealing with very large datasets of large diagrams, since all the previous works dealing with this data [LY18, CCO17, AEK⁺17] use kernel methods. Results are displayed in Table 6.1. As it can be observed, we improve on previous results for ORBIT5K, we also show with ORBIT100K that classification accuracy is further increased as more observations are made available. For consistency we use the same accuracy metric as [LY18], that is, we split observations in 70%-30% training-test sets and report the average test accuracy over 100 runs.

Graph classification. The second experiment we used to illustrate the use of PERSLAY as an interface to incorporate persistence diagrams in learning pipeline is a *graph classification* task. Each graphs in a given dataset was turned into a set of persistence diagrams using the so-called *Heat Kernel Signatures* (HKS) of the graph. For the sake of concision, we do not give a detailed presentation of the HKS signatures here, and refer the reader to the original paper for complementary information.

¹http://gudhi.gforge.inria.fr/python/latest/alpha_complex_ref.html

We used a series of different graph datasets commonly used as a baseline in graph classification problems.

- REDDIT5K, REDDIT12K, COLLAB (from [YV15]) IMDB-B, IMDB-M (from [TVH18]) are composed of social graphs.
- COX2, DHFR, MUTAG, PROTEINS, NCI1, NCI109 are graphs coming from medical or biological frameworks (also from [TVH18]).

We compared performances with five other top graph classification methods.

- Scale-variant topo [TVH18] leverages a kernel for ordinary persistence diagrams computed on point cloud used to encode the graphs.
- RetGK [ZWX⁺18] is a kernel method for graphs that leverages eventual *attributes* on the graph vertices and edges.
- FGSD [VZ17] is a finite-dimensional graph embedding that does not leverage attributes.
- Finally, GCNN [XC19] and GIN [XHLJ19] are two graph neural network approaches that reach top-tier results.

One could also compare our results on the REDDIT datasets to the ones of [HKNU17], where authors also use persistence diagrams to feed a network (using as first channel a particular case of PERSLAY), achieving 54.5% and 44.5% of accuracy on REDDIT5K and REDDIT12K respectively.

In this experiment, we used a voluntarily simple network architecture, namely a two-layer network. The first layer is PERSLAY, which processes persistence diagrams. The resulting vector is normalized and fed to the second and final layer, a fully-connected layer whose output is used for predictions. See Figure 6.3 for an illustration. We emphasize that this simplistic two-layer architecture is designed so as to produce understanding rather than achieving the best possible performances.




Figure 6.3: Network architecture illustrated in the case of our graph classification experiments. Each graph is encoded as a set of 4 persistence diagrams, then processed by an independent instance of PERSLAY. Each instance embeds diagrams in some vector space using two functions w, ϕ that are optimized during training and a fixed permutation-invariant operator op .




Figure 6.4: Weight function w when chosen to be a grid with size 20×20 before and after training (MUTAG dataset). Here, Ord0, Rel1, Ext0, and Ext1 denote the different diagrams extracted from a graphs, each of them being processed by an independent instance of PERSLAY.

Dataset	SV ¹	RetGK* ²	FGSD ³	GCNN ⁴	GIN ⁵	PERSLAY	
						Mean	Max
REDDIT5K	—	56.1	47.8	52.9	57.0	55.6	56.5
REDDIT12K	—	48.7	—	46.6	—	47.7	49.1
COLLAB	—	81.0	80.0	79.6	80.1	76.4	78.0
IMDB-B	72.9	71.9	73.6	73.1	74.3	71.2	72.6
IMDB-M	50.3	47.7	52.4	50.3	52.1	48.8	52.2
COX2*	78.4	80.1	—	—	—	80.9	81.6
DHFR*	78.4	81.5	—	—	—	80.3	80.9
MUTAG*	88.3	90.3	92.1	86.7	89.0	89.8	91.5
PROTEINS*	72.6	75.8	73.4	76.3	75.9	74.8	75.9
NCI1*	71.6	84.5	79.8	78.4	82.7	73.5	74.0
NCI109*	70.5	—	78.8	—	—	69.5	70.1

Table 6.2: Classification accuracy over benchmark graph datasets. Our results (PersLay, right hand side) are recorded from ten runs of a 10-fold classification evaluation. “Mean” is consistent with [ZWX+18]², while “Max” should be compared to [TVH18]¹, [VZ17]³, [XC19]⁴ and [XHLJ19]⁵, as it corresponds to the mean accuracy over a *single 10-fold*. The * indicates datasets that contain attributes (labels) on graph nodes and symmetrically the methods that leverage such attributes for classification purposes.

CHAPTER 7

COMPLEMENTARY EXAMPLES

Abstract

This final chapter gathers three independent examples that showcase how adopting an optimal transport perspective and a measure-based formalism is helpful when dealing with persistence diagrams. Section 7.1 is dedicated to the study of *expected persistence diagrams*, that arise when considering persistence diagrams in a random setting. As those are measures with continuous support, they can be manipulated using our formalism, where standard definitions of diagram metrics would be insufficient. Section 7.2 studies the *quantization* of persistence diagrams, namely approximating a diagram by a persistence measure with a chosen size support. Finally, Section 7.3 addresses the question of estimating *shift-invariant* distances between persistence diagrams. Using the *entropic regularization* of optimal transport (adapted to persistence diagrams), we obtain a simple and fast algorithm (although non-convex) to propose solutions to this problem.

Implementation resources.

- [A tutorial for Expected persistence diagrams.](#)
- (To appear in [Gudhi](#)) Quantization of persistence diagrams.
- (In progress) Shift-invariant distance.

7.1 Expected persistence diagrams and their estimators

Recall that \mathcal{M}^p denotes the space of persistence measures (Radon measures supported on the open half-plane Ω) for some parameter $1 \leq p < \infty$, and $\mathcal{D}^p \subset \mathcal{M}^p$ denotes the subspace of persistence diagrams. These spaces are equipped with metric OT_p . These are Polish metric spaces (see Chapter 3) and therefore, one can consider probability distribution supported on these.

When observing a sample of persistence diagrams $\mu_1 \dots \mu_K \in \mathcal{D}^p$, a natural model is to assume that the μ_i s are sampled i.i.d. with respect to some underlying probability distribution P belonging to $\mathcal{W}^p(\mathcal{D}^p)$, the Wasserstein space on \mathcal{D}^p (see Section 2.2 and Chapter 3 for details). This section is dedicated to the study of a natural statistical descriptor that one can build on P : its (linear) expectation, denoted by $\mathbb{E}_P[\boldsymbol{\mu}]$ where $\boldsymbol{\mu} \sim P$, and called its *expected persistence diagram* (EPD). We start by proving a stability result on the map $P \mapsto \mathbb{E}_P[\boldsymbol{\mu}]$ (this result was part of [DL19, §5]). As in practice one only has access to a finite sample $\mu_1 \dots \mu_K \sim P$, we will study how the empirical estimator $\frac{1}{K} \sum \mu_i$ performs when it comes to model the underlying object $\mathbb{E}_P[\boldsymbol{\mu}]$.

7.1.1 Definition and properties

Let P be a probability distribution of persistence diagrams. One may want to build a statistic on P , such as defining and computing its mean. As $(\mathcal{D}^p, \text{OT}_p)$ is not a vector space but only a metric space, the notion of mean is ill-defined. A natural attempt is to consider *Fréchet means*, that is minimizers of the functional $\nu \mapsto \int_{\mu \in \mathcal{D}^p} \text{OT}_p^p(\nu, \mu) dP(\mu)$. Algorithms to address this minimization problem have been proposed (see Chapter 5 and references therein) but those only provide an estimate and can be expensive to compute. On the other hand, one could go for a much simpler approach and define a measure $\mathbb{E}_P[\boldsymbol{\mu}]$ supported on Ω by stating, for $A \subset \Omega$ compact,

$$\mathbb{E}_P[\boldsymbol{\mu}](A) := \mathbb{E}_P[\boldsymbol{\mu}(A)], \quad (7.1)$$

where $\boldsymbol{\mu} \sim P$, and $\boldsymbol{\mu}(A)$ is the (random) number of points of $\boldsymbol{\mu}$ that belongs to A . This deterministic measure, called the *expected persistence diagram* of P , was introduced in [DC19]. In this work, Divol and Chazal proved in particular that under mild assumptions on P , the corresponding EPD admits a density with respect to the Lebesgue measure on Ω . Therefore, it is an element of \mathcal{M}^p , although not being a persistence diagram.




Figure 7.1: (Left) A n -sample on a torus, with $n = 500$. (Middle) The corresponding persistence diagram. (Right) The corresponding (empirical) expected persistence diagram, estimated by sampling $K = 100$ times n points on a torus. Note that for illustration purpose, the empirical EPD has been rendered using a 30×30 grid, and masses of each bin have been weighted according to their distance to the diagonal, to the power $p = 2$.

Intuitively, the mass $\mathbb{E}_P[\boldsymbol{\mu}]$ gives to an area $A \subset \Omega$ is the average number of points that fall in A when sampling diagrams with respect to P . See Figure 7.1 for an illustration.

REMARK 7.1. A well-suited definition of the linear expectation in (7.1) requires technical care (basically, turning the finite sum into a Bochner integral) and is detailed in Section 7.1.3. This is required, in particular, to apply the Jensen inequality in our proof of the stability result. Nonetheless, on first read one can treat $\mathbb{E}_P[\boldsymbol{\mu}]$ as a standard expectation and take Equation (7.1) as a definition.

The main theoretical result of this section states the stability of the map

$$\begin{aligned} \mathcal{W}^p(\mathcal{D}^p) &\rightarrow \mathcal{M}^p \\ P &\mapsto \mathbb{E}_P[\boldsymbol{\mu}]. \end{aligned} \quad (7.2)$$

PROPOSITION 7.2. Let $P, P' \in \mathcal{W}^p(\mathcal{M}^p)$. We have

$$\text{OT}_p(\mathbb{E}_P[\boldsymbol{\mu}], \mathbb{E}_{P'}[\boldsymbol{\mu}]) \leq W_{p, \text{OT}_p}(P, P').$$

The proof can be found in Section 7.1.3.

Using stability results on the distances $\mathbf{d}_p (= \text{OT}_p)$ between persistence diagrams (see Theorem 2.7 or [CSEHM10]), one is able to obtain a more precise control between the expectations in some situations. For \mathbb{Y} a sample




Figure 7.2: Just as standard diagrams, expected persistence diagrams are stable with respect to small perturbations. (Left) Two point clouds (with $n = 130$ points). Blue (dot) is (randomly, iid) sampled exactly on the union of two circles, while red (cross) is sampled on the union of these two circles in addition to some small bounded noise. (Middle and Right) The (empirical) expected persistence diagram (EPD) for the “clean” (blue) sampling process and for the “noisy” (red) sampling process respectively. Due to the noise, the second EPD has some mass close to the diagonal; but the two EPDs remain close in terms of OT_∞ metric.

in some metric space, denote by $\text{Dgm}(\mathbb{Y})$ the persistence diagram of \mathbb{Y} built with the Čech filtration.

PROPOSITION 7.3. *Let ξ, ξ' be two probability measures on \mathbb{R}^d . Let \mathbb{X}_n (resp. \mathbb{X}'_n) be a n -sample of law ξ (resp. ξ'). Then, for any $k > d$, and any $p \geq k + 1$,*

$$\text{OT}_p^p(\mathbb{E}[\text{Dgm}(\mathbb{X}_n)], \mathbb{E}[\text{Dgm}(\mathbb{X}'_n)]) \leq C_{k,d} \cdot n \cdot W_{p-k}^{p-k}(\xi, \xi') \quad (7.3)$$

where $C_{k,d} := C \text{diam}(\mathbb{X})^{k-d} \frac{k}{k-d}$ for some constant C depending only on \mathbb{X} . In particular, letting $p \rightarrow \infty$, we obtain a bottleneck stability result:

$$\text{OT}_\infty(\mathbb{E}[\text{Dgm}(\mathbb{X}_n)], \mathbb{E}[\text{Dgm}(\mathbb{X}'_n)]) \leq W_\infty(\xi, \xi'). \quad (7.4)$$

Proof. Let γ be any coupling between \mathbb{X}_n a n -sample of law ξ , and \mathbb{X}'_n a n -sample of law ξ' . According to Proposition 7.2,

$$\text{OT}_p^p(\mathbb{E}[\text{Dgm}(\mathbb{X}_n)], \mathbb{E}[\text{Dgm}(\mathbb{X}'_n)]) \leq \mathbb{E}_\gamma[\text{OT}_p^p(\text{Dgm}(\mathbb{X}_n), \text{Dgm}(\mathbb{X}'_n))].$$

Theorem 2.7 states (see also [CSEHM10, Wasserstein Stability Theorem]) that

$$\text{OT}_p^p(\text{Dgm}(\mathbb{X}_n), \text{Dgm}(\mathbb{X}'_n)) \leq C_{k,d} H(\mathbb{X}_n, \mathbb{X}'_n)^{p-k},$$

where $C_{k,d} := C \text{diam}(\mathbb{X})^{k-d} \frac{k}{k-d}$ for some constant C depending only on \mathbb{X} , and H is the Hausdorff distance between sets. By taking the infimum on transport plans γ , we obtain

$$\text{OT}_p^p(\mathbb{E}[\text{Dgm}(\mathbb{X}_n)], \mathbb{E}[\text{Dgm}(\mathbb{X}'_n)]) \leq C_{k,d} W_{H,p-k}^{p-k}(\xi^{\otimes n}, (\xi')^{\otimes n}),$$

where $W_{H,p}$ is the p -Wasserstein distance between probability distributions on compact sets of the manifold \mathbb{X} , endowed with the Hausdorff distance. Lemma 15 of [CFL⁺15] states that

$$W_{H,p-k}^{p-k}(\xi^{\otimes n}, (\xi')^{\otimes n}) \leq n \cdot W_{p-k}^{p-k}(\xi, \xi'),$$

concluding the proof. \square

Note that this proposition illustrates the usefulness of introducing new distances OT_p : considering the proximity between linear expectations requires to extend the metrics \mathbf{d}_p to Radon measures.

7.1.2 Approximating EPDs with empirical EPDs.

As in practice, one generally only have access to a sample of observations instead of the whole underlying law P and thus the corresponding expected persistence diagram \mathbb{E}_P , we now address the question of estimating the EPD by its empirical counterpart.

Let P be a probability measure on $(\mathcal{D}_p, \text{OT}_p)$ and let $\boldsymbol{\mu}^{(1)}, \dots, \boldsymbol{\mu}^{(K)}$ be a K -sample¹ of law P . The empirical EPD is defined as

$$\bar{\boldsymbol{\mu}}_K := \frac{\boldsymbol{\mu}^{(1)} + \dots + \boldsymbol{\mu}^{(K)}}{K}. \quad (7.5)$$

We first study the convergence of $\bar{\boldsymbol{\mu}}_K$ to $\mathbb{E}_P[\boldsymbol{\mu}]$ as $K \rightarrow \infty$.

PROPOSITION 7.4. *Let $p \geq 1$ and assume that $\mathbb{E}_P[\text{Pers}_p(\boldsymbol{\mu})] < \infty$. Then, $\mathbb{E}[\text{OT}_p^p(\bar{\boldsymbol{\mu}}_K, \mathbb{E}_P[\boldsymbol{\mu}])] \rightarrow 0$.*

Proof. We use the characterization of convergence that we proved in Theorem 3.7. It states that the convergence $\text{OT}_p^p(\bar{\boldsymbol{\mu}}_K, \mathbb{E}_P[\boldsymbol{\mu}]) \rightarrow 0$ holds almost surely if and only if (i) $\text{Pers}_p(\bar{\boldsymbol{\mu}}_K) \rightarrow \text{Pers}_p(\mathbb{E}_P[\boldsymbol{\mu}])$ a.s. and (ii) $\bar{\boldsymbol{\mu}}_K$ converges vaguely to $\mathbb{E}_P[\boldsymbol{\mu}]$ a.s., that is if $\bar{\boldsymbol{\mu}}_K(\phi) \rightarrow \mathbb{E}_P[\boldsymbol{\mu}](\phi)$ a.s., for all $\phi : \Omega \rightarrow \mathbb{R}$ continuous with compact support [Kal17, Lemma 4.8]. Those two conditions hold by the strong law of large numbers. Thus, we have

¹That is, K measures sampled i.i.d.

$\text{OT}_p^p(\bar{\boldsymbol{\mu}}_{\mathbf{K}}, \mathbb{E}_P[\boldsymbol{\mu}]) \rightarrow 0$ a.s. To obtain the convergence of the expectation, we also need the sequence of random variables $(\text{OT}_p^p(\bar{\boldsymbol{\mu}}_{\mathbf{K}}, \mathbb{E}_P[\boldsymbol{\mu}]))_K$ to be uniformly integrable. By the triangle inequality, and using that $(a+b)^p \leq 2^{p-1}(a^p + b^p)$ for $a, b \geq 0$, with $\mathbf{0}$ denoting the empty diagram,

$$\begin{aligned} \text{OT}_p^p(\bar{\boldsymbol{\mu}}_{\mathbf{K}}, \mathbb{E}_P[\boldsymbol{\mu}]) &\leq (\text{OT}_p(\bar{\boldsymbol{\mu}}_{\mathbf{K}}, \mathbf{0}) + \text{OT}_p(\mathbb{E}_P[\boldsymbol{\mu}], \mathbf{0}))^p \\ &\leq 2^{p-1}(\text{Pers}_p(\bar{\boldsymbol{\mu}}_{\mathbf{K}}) + \text{Pers}_p(\mathbb{E}_P[\boldsymbol{\mu}])) \\ &= 2^{p-1} \frac{1}{K} \sum_{k=1}^K \text{Pers}_p(\mu^{(\mathbf{k})}) + \mathbb{E}_P[\text{Pers}_p(\boldsymbol{\mu})]. \end{aligned}$$

This last quantity is uniformly integrable, and we can conclude as almost sure convergence and uniform integrability imply convergence in expectation. \square

We can make this convergence result more quantitative by making more assumptions on the probability P , namely that the diagrams sampled according to P have a uniformly bounded mass (number of points) and support.

PROPOSITION 7.5. *Consider $p \geq 1$ and let P be a probability measure on $(\mathcal{D}_p, \text{OT}_p)$ such that $\boldsymbol{\mu} \sim P$ is a.s. supported on $[-L, L]^2$ and of total mass less than or equal to M . Then, for any $K > 0$,*

$$\mathbb{E}[\text{OT}_p^p(\bar{\boldsymbol{\mu}}_{\mathbf{K}}, \mathbb{E}_P[\boldsymbol{\mu}])] \leq \frac{C_p M L^p}{\sqrt{K}} a_p(K). \quad (7.6)$$

with $a_p(K) = 1$ if $p > 1$ and $a_1(K) = \ln(K)$.

The proof, slightly technical, is delayed to Section 7.1.3.

REMARK 7.6. *Using the dual formulation (Section 3.4) for $p = 1$, the quantity $\text{OT}_1(\mu, \nu)$ is equal to $\sup_f |\mu(f) - \nu(f)|$, where the supremum is taken on all Lipschitz continuous functions $f : \bar{\Omega} \rightarrow \mathbb{R}$ with $f(\partial\Omega) = 0$. Hence, Proposition 7.5 implies that*

$$\mathbb{E}[\sup_f |\bar{\boldsymbol{\mu}}_{\mathbf{K}}(f) - \mathbb{E}_P[\boldsymbol{\mu}](f)|] \leq C M L \frac{\ln(K)}{\sqrt{K}}. \quad (7.7)$$

Taking different choices of f induce different representations on PDs, including e.g. persistent surfaces [AEK⁺17], persistent silhouettes [CFL⁺14], the persistence scale-space kernel [RHBK15] or gaussian functions centered at random locations [HKNU17]. Hence, Proposition 7.5 gives a uniform rate of convergence for all those representations.

Let $\mathcal{M}_{M,L}$ be the set of persistence measures with support included in $[-L, L]^2$ and of mass less than or equal to M , and let $\mathcal{P}_{M,L}$ be the set of all probability distributions supported on the set $\mathcal{D} \cap \mathcal{M}_{M,L}$. The following proposition shows that the minimax rate for the estimation of the EPD given K observations of law $P \in \mathcal{P}_{M,L}$ is of order $1/\sqrt{K}$, implying that the empirical EPD is a minimax estimator of the EPD. The notation $a_K \asymp b$ means that there exists a constant C such that $\forall K, a_K/b \leq C$ and $b/a_K \leq C$.

PROPOSITION 7.7. *Let $1 \leq p < \infty$. One has*

$$\inf_{\hat{\mu}_{\mathbf{K}}} \sup_{P \in \mathcal{P}_{M,L}} \mathbb{E}[\text{OT}_p^p(\hat{\mu}_{\mathbf{K}}, \mathbb{E}_P[\mu])] \asymp \frac{ML^p}{\sqrt{K}}, \quad (7.8)$$

where the infimum is taken on all measurable $\mathcal{M}_{M,L}$ -valued functions $\hat{\mu}_{\mathbf{K}}$ of a K -sample $\mu^{(1)}, \dots, \mu^{(K)}$ of law $P \in \mathcal{P}_{M,L}$.

7.1.3 Complementary details and delayed proofs

Formal definition of the expected persistence diagram. Define \mathcal{M}_{\pm} the space of signed measures on Ω , i.e. a measure $\mu \in \mathcal{M}_{\pm}$ is written $\mu_+ - \mu_-$ for two finite (non-negative) measures $\mu_+, \mu_- \in \mathcal{M}_f$. The total variation distance $|\cdot|$ is a norm on \mathcal{M}_{\pm} , and $(\mathcal{M}_{\pm}, |\cdot|)$ is a Banach space. The Bochner integral [Boc33] is a generalization of the Lebesgue integral for functions taking their values in Banach space. We define the expected persistence measure of $P \in \mathcal{W}^p(\mathcal{M}^p)$ as the Bochner integral of some pushforward of P . More precisely, recall the definition (3.8) of $\mu^{[p]}$ and define

$$F : (\mathcal{M}^p, \text{OT}_p) \rightarrow (\mathcal{M}_{\pm}, |\cdot|) \\ \mu \mapsto \mu^{[p]}.$$

Note that F has an inverse G on \mathcal{M}_f , defined by $G(\nu)(A) := \int_A d(x, \Omega)^{-p} d\nu(x)$ for $A \subset \Omega$ a Borel set. Theorem 3.7 implies that G is a continuous function from $(\mathcal{M}_f, |\cdot|)$ to $(\mathcal{M}^p, \text{OT}_p)$. In particular, as \mathcal{M}_f and \mathcal{M}^p are Polish spaces and G is injective, F is measurable (see [Kec95, Theorem 15.1]). For $P \in \mathcal{W}^p(\mathcal{M}^p(\Omega))$, define for $\mu \sim P$, $\mathbb{E}[\mu]$ the linear expectation of P by

$$\mathbb{E}[\mu] := G \left(\int \nu d(F_{\#}P)(\nu) \right) \in \mathcal{M}^p, \quad (7.9)$$

where the integral is the Bochner integral on the Banach space $(\mathcal{M}_{\pm}, |\cdot|)$ and $F_{\#}P$ is the pushforward of P by F . One can check that $\mathbb{E}[\mu]$ defined in that way satisfies the relation

$$\forall K \subset \Omega \text{ compact}, \quad \mathbb{E}[\mu](K) = \mathbb{E}[\mu(K)].$$

Delayed proofs.

Proof of Proposition 7.2. The proof consists in applying Jensen's inequality in an infinite-dimensional setting. We first recall that the function OT_p^p is convex (Lemma 3.13). We then use the following result, which is a particular case of [Per74, Theorem 3.10].

PROPOSITION 7.8. *Let \mathcal{X} be a Hausdorff locally convex topological vector space and $C \subset \mathcal{X}$ a closed convex set. Let Q be a probability measure on \mathcal{X} endowed with its borelian σ -algebra, which is supported on C . Assume that $\int \|x\| dQ(x) < \infty$. Let $f : C \rightarrow [0, \infty)$ be a continuous convex function with $\int f(x) dQ(x) < \infty$. Then*

$$f\left(\int x dQ(x)\right) \leq \int f(x) dQ(x).$$

Let $\mathcal{X} = \mathcal{M}_\pm \times \mathcal{M}_\pm$ which is a Banach space (endowed with the product norm), and thus in particular a Hausdorff locally convex topological vector space. Let $C = \mathcal{M}_f \times \mathcal{M}_f$, which is convex and closed (closedness follows immediately from the definition of the total variation $|\cdot|$) and let $f = \text{OT}_p^p \circ (G, G) : \mathcal{X} \rightarrow \mathbb{R}$. The continuity of G implies that f is continuous and Lemma 3.13 implies the convexity of f . Let P, P' be two probability measures in $\mathcal{W}^p(\mathcal{M}^p)$ and γ be an optimal coupling between P and P' . We let Q be the image measure of γ by (F, F) , so that

$$\begin{aligned} \int_{x \in \mathcal{X}} \|x\| dQ(x) &= \int_{\mu, \mu' \in \mathcal{M}^p} \max(|\mu^{(p)}|, |(\mu')^{(p)}|) d\gamma(\mu, \mu') \\ &\leq \int_{\mu} \text{Pers}_p(\mu) dP(\mu) + \int_{\mu'} \text{Pers}_p(\mu') dP'(\mu') < \infty \end{aligned}$$

and that

$$\int_{x \in \mathcal{X}} f(x) dQ(x) = \int_{\mu, \mu' \in \mathcal{M}^p} \text{OT}_p^p(\mu, \mu') d\gamma(\mu, \mu') = W_{p, \text{OT}_p}^p(P, P') < \infty.$$

Also, we have

$$\begin{aligned} \int x dQ(x) &= \int_{\nu, \nu' \in \mathcal{M}^p} (\nu, \nu') d(F, F)_\# \gamma(\nu, \nu') \\ &= \left(\int_{\nu \in \mathcal{M}^p} \nu dF_\# P(\nu), \int_{\nu' \in \mathcal{M}^p} \nu' dF_\# P'(\nu') \right), \end{aligned}$$

so that by (7.9), $f(\int x dQ(x)) = \text{OT}_p^p(\mathbb{E}[\boldsymbol{\mu}], \mathbb{E}[\boldsymbol{\mu}'])$, where $\boldsymbol{\mu} \sim P$ and $\boldsymbol{\mu}' \sim P'$. Proposition 7.2 yields the conclusion. \square

Proof of Proposition 7.5. The total variation between measures will be denoted by $|\cdot|$ in the proof. By using the triangle inequality,

$$\begin{aligned} \text{OT}_p(\mu, \nu) &\leq \text{OT}_p\left(\frac{|\nu|}{|\mu|}\mu, \mu\right) + \text{OT}_p\left(\frac{|\nu|}{|\mu|}\mu, \nu\right) \\ &\leq |\mu - \nu|^{1/p} L + W_{p, \|\cdot\|}\left(\frac{|\nu|}{|\mu|}\mu, \nu\right), \end{aligned}$$

where at the second line, the first term is bounded by considering the transport plan which consists in taking the identity map on Ω , and by mapping the remaining mass on the diagonal, with a cost therefore bounded by $|\mu - \nu| L^p$; and the second term is bounded by observing that any off-diagonal transport plan π supported on $\Omega \times \Omega$ between $\frac{|\nu|}{|\mu|}\mu$ and ν (which do have the same total mass) is an admissible (partial) plan between these two persistence measures.

Let \mathcal{P}_l be the natural partition of $[-L, L]^2$ induced by a $2^l \times 2^l$ regular grid. For any μ, ν supported on $[-L, L]^2$, we let

$$D_p(\mu, \nu) := \sum_{l \geq 1} 2^{-pl} \sum_{F \in \mathcal{P}_l} |\mu(F) - \nu(F)|. \quad (7.10)$$

This quantity was introduced by [FG15], who proved in particular that if two measures σ, τ have the same total mass, then there exists a constant κ_p such that $W_{p, \|\cdot\|}^p(\sigma, \tau) \leq \kappa_p L^p D_p(\sigma, \tau)$. Also, for $F \in \mathcal{P}_l$, we have $\left|\frac{|\nu|}{|\mu|}\mu(F) - \nu(F)\right| \leq \left|\frac{|\nu|}{|\mu|} - 1\right| \mu(F) + |\mu(F) - \nu(F)|$. Hence,

$$\begin{aligned} W_{p, \|\cdot\|}\left(\frac{|\nu|}{|\mu|}\mu, \nu\right) &\leq \kappa_p L^p D_p\left(\frac{|\nu|}{|\mu|}\mu, \nu\right) \\ &\leq \kappa_p L^p \sum_{l \geq 1} 4^{-l} \sum_{F \in \mathcal{P}_l} \left|\frac{|\nu|}{|\mu|} - 1\right| \mu(F) + \kappa_p L^p D_p(\mu, \nu) \\ &\leq \frac{\kappa_p}{3} L^p |\mu - \nu| + \kappa_p L^p D_p(\mu, \nu), \end{aligned}$$

and using that $(a + b)^p \leq 2^{p-1}(a^p + b^p)$, we obtain that $\text{OT}_p^p(\mu, \nu) \leq L^p \left(2^{p-1} + \frac{\kappa_p}{3}\right) |\mu - \nu| + 2^{p-1} \kappa_p L^p D_p(\mu, \nu)$. We now apply this last inequality to $\mu = \mathbb{E}_P[\boldsymbol{\mu}]$ and $\nu = \bar{\boldsymbol{\mu}}_{\mathbf{K}}$. For $A \subset [-1, 1]^2$ a measurable set, we have

$$\mathbb{E}[|\bar{\boldsymbol{\mu}}_{\mathbf{K}}(A) - \mathbb{E}_P[\boldsymbol{\mu}](A)|] \leq \sqrt{\frac{\mathbb{E}_P[\boldsymbol{\mu}(A)^2]}{K}} \leq \frac{M}{\sqrt{K}}. \quad (7.11)$$

Thus, $\mathbb{E}[|\bar{\boldsymbol{\mu}}_{\mathbf{K}} - \mathbb{E}_P[\boldsymbol{\mu}]|] \leq M/\sqrt{K}$ and, by using Jensen inequality twice,

$$\begin{aligned} & \sum_{F \in \mathcal{P}_l} \mathbb{E}_P[|\bar{\boldsymbol{\mu}}_{\mathbf{K}}(F) - \mathbb{E}_P[\boldsymbol{\mu}](F)|] \\ & \leq 2^l \left(\sum_{F \in \mathcal{P}_l} \mathbb{E}_P[(\bar{\boldsymbol{\mu}}_{\mathbf{K}}(F) - \mathbb{E}_P[\boldsymbol{\mu}](F))^2] \right)^{1/2} \\ & \leq \frac{2^l}{\sqrt{K}} \left(\sum_{F \in \mathcal{P}_l} \mathbb{E}_P[\boldsymbol{\mu}(F)^2] \right)^{1/2} \\ & \leq \frac{2^l}{\sqrt{K}} \left(M \sum_{F \in \mathcal{P}_l} \mathbb{E}_P[\boldsymbol{\mu}(F)] \right)^{1/2} \leq \frac{2^l M}{\sqrt{K}}. \end{aligned}$$

In the case $p > 1$, we obtain the conclusion as $\sum_{l \geq 1} 2^{(1-p)l} = c_p < \infty$. In the case $p = 1$, we use that

$$\sum_{F \in \mathcal{P}_l} \mathbb{E}_P[|\bar{\boldsymbol{\mu}}_{\mathbf{K}}(F) - \mathbb{E}_P[\boldsymbol{\mu}](F)|] \leq 2 \sum_{F \in \mathcal{P}_l} \mathbb{E}_P[\boldsymbol{\mu}(F)] \leq 2M,$$

yielding

$$\begin{aligned} \mathbb{E}[D_1(\bar{\boldsymbol{\mu}}_{\mathbf{K}}, \mathbb{E}_P[\boldsymbol{\mu}])] & \leq \sum_{l \geq 1} 2^{-l} \min \left(2M, \frac{2^l M}{\sqrt{K}} \right) \\ & \leq CM \frac{\ln(K)}{\sqrt{K}} \end{aligned}$$

for some absolute constant $C > 0$. □

Proof of Proposition 7.7. The upper bound on the minimax rate is given by the rates of convergence of the empirical EPD given in Proposition 7.5.

To obtain the lower bound, we use that if μ, ν are two persistence measures on Ω of mass less than or equal to M , then we can use the results proved in Section 3.2, namely Proposition 3.15 which states

$$\text{OT}_p(\mu, \nu) = W_{p, \rho}(\Phi(\mu), \Phi(\nu)),$$

where

- ρ is a distance on $\tilde{\Omega} := \Omega \cup \{\partial\Omega\}$ defined by $\forall x, y \in \tilde{\Omega}$,

$$\rho(x, y) = \min(d(x, y), d(x, \partial\Omega) + d(y, \partial\Omega)),$$

where by convention for $x \in \Omega$, we set $d(x, \partial\Omega) := d(x, \mathbf{p}_{\partial\Omega}(x))$, with $\mathbf{p}_{\partial\Omega}(x)$ the orthogonal projection of x onto the diagonal $\partial\Omega$.

- $\Phi(\mu) := \mu + (2M - |\mu|)\delta_{\partial\Omega}$.
- $W_{p,\rho}$ denotes the Wasserstein distance between measures (with same mass) with distance ρ (see Section 2.2.2).

As Φ is a bijection, the minimax rates for the estimation of $\mathbb{E}_P[\mu]$ is therefore equal to

$$\inf_{\Phi(\hat{\mu}_{\mathbf{K}})} \sup_{P \in \mathcal{P}_{M,L}} \mathbb{E}[W_{p,\rho}^p(\Phi(\hat{\mu}_{\mathbf{K}}), \Phi(\mathbb{E}_P[\mu]))].$$

The set $\mathcal{P}_{M,L}$ contains in particular the set of all distributions P_τ for which $\mu \sim P$ satisfies $\Phi(\mu) = 2M\delta_{\mathbf{x}}$, where $\mathbf{x} \sim \tau$, where τ is a probability measure on $(\tilde{\Omega}, \rho)$. For such a distribution P , one has $\Phi(\mathbb{E}_P[\mu]) = 2M\tau$, so that, letting \mathcal{Q} be the set of Borel probability measures on $(\tilde{\Omega} \cap [-L, L]^2, \rho)$, the minimax rate must be larger than

$$\inf_{\hat{\mathbf{a}}_{\mathbf{K}}} \sup_{\tau \in \mathcal{Q}} \mathbb{E}[W_{p,\rho}^p(\hat{\mathbf{a}}_{\mathbf{K}}, 2M\tau)],$$

where the infimum is taken on all measurable functions based on K observations of the form $2M\delta_{\mathbf{x}^{(k)}}$ with $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(K)}$ a K -sample of law $\tau \in \mathcal{Q}$. Hence, we have shown that the minimax rate for the estimation of $\mathbb{E}_P[\mu]$ with respect to OT_p is larger up to a factor M than the minimax rate for the estimation of τ a distribution on $(\tilde{\Omega} \cap [-L, L]^2, \rho)$ given K i.i.d. observations of τ . As the minimax rate for this problem is known to be larger than L^p/\sqrt{K} [SP18, Theorem 9], we obtain the conclusion. \square

7.2 Quantization of persistence diagrams

We address in this section the *quantization* of persistence diagram. By quantization, we mean a *projection* with respect to diagram metrics on a subspace of the space of persistence diagrams (or, as we will see, general persistence measures) consisting of measures with a fixed (finite) support size. Let \mathcal{M}^p denote the space of persistence measures and $\mathcal{D}^p \subset \mathcal{M}^p$ be the space of persistence diagrams, for some parameter $1 \leq p \leq \infty$. Recall that \mathcal{M}^p is equipped with a metric OT_p which generalizes the partial matching metric \mathbf{d}_p defined on \mathcal{D}^p , see Section 2.1 for details. As we target numerical applications, we implicitly only consider in this section persistence diagrams and measures with *finite* support.

7.2.1 Definition and motivations

DEFINITION 7.9. Fix some integer $k \geq 1$. Let $\mathcal{M}_k \subset \mathcal{M}^p(\overline{\Omega})$ be the persistence measures with support of size k supported on $\overline{\Omega}$, that is

$$\mathcal{M}_k := \left\{ \mu(a, X) := \sum_{i=1}^k a_i \delta_{x_i}, \ a \in \mathbb{R}_+^k, X = (x_1 \dots x_k) \in \overline{\Omega}^k \right\}. \quad (7.12)$$

A k -quantization is a projection onto a subset of \mathcal{X} of \mathcal{M}_k , that is a map

$$Q : \nu \mapsto Q(\nu) \in \arg \min_{\mu \in \mathcal{X}} \{\text{OT}_p^p(\mu, \nu)\}, \quad (7.13)$$

In the vocabulary of quantization, the x_k s are called *codepoints* and $X = (x_1 \dots x_k)$ is called a *codebook*.

REMARK 7.10. Note that quantized diagrams are supported on $\overline{\Omega}$ (i.e. are allowed to have points on the diagonal). This gives, for instance, a solution to the problem $\min_x \text{OT}_p(\delta_x, \mathbf{0})$, where δ_x is a persistence diagram with a single point $x \in \overline{\Omega}$ and $\mathbf{0}$ represents the empty diagram.

As it will be used in the following, we recall that metrics OT_p can be computed as solution of a linear program (as introduced in Chapter 5). Let us consider two persistence measures with finite support $\mu = \sum_{i=1}^{n_1} a_i \delta_{x_i}$ and $\nu = \sum_{j=1}^{n_2} b_j \delta_{y_j}$, where $a_i, b_j \geq 0$ and $x_i, y_j \in \overline{\Omega}$. Let $X = (x_i)_i, Y = (y_j)_j$ denote the locations of the points and $a = (a_i)_i, b = (b_j)_j$ be the corresponding masses. Let $\mathbf{p}_{\partial\Omega}(x)$ denote the (orthogonal) projection of $x \in \overline{\Omega}$ onto the diagonal $\partial\Omega := \{(t, t), t \in \mathbb{R}\}$. Let also $\text{pers}(x)$ denote the distance of $x \in \overline{\Omega}$ to the diagonal, that is $\text{pers}(x) := d(x, \mathbf{p}_{\partial\Omega}(x))$.

Let $C = C(X, Y)$ be the cost matrix with block structure

$$C = \begin{pmatrix} C_{XY} & C_{X\Delta} \\ C_{\Delta Y} & 0 \end{pmatrix} \in \mathbb{R}^{(n_1+1) \times (n_2+1)}, \quad (7.14)$$

with $C_{X\Delta} = (\text{pers}(x_i)^p)_i$, $C_{\Delta Y} = (\text{pers}(y_j)^p)_j$ and $C_{XY} = (d(x_i, y_j)^p)_{ij}$.

It reads

$$\text{OT}_p^p(\mu, \nu) = \mathbf{L}_C(a, b) \quad (7.15)$$

where \mathbf{L}_C is defined by the following optimization problems, dual of each other:

$$\mathbf{L}_C(a, b) = \min_{P \in \Pi(a, b)} \langle C, P \rangle = \min_{P \in \Pi(a, b)} \sum_{ij} C_{ij} P_{ij}, \quad (7.16)$$

$$= \max_{f, g \in \Psi_C} \sum_{i=1}^{n_1} a_i f_i + f_\Delta \left(\sum_j b_j \right) + \sum_{j=1}^{n_2} b_j g_j + g_\Delta \left(\sum_i a_i \right), \quad (7.17)$$

where f_Δ and g_Δ respectively denote the $(n_1 + 1)$ -th and $(n_2 + 1)$ -th coordinates of f and g , and the constraint sets are respectively defined as

$$\begin{aligned} \Pi(a, b) &:= \{P \in \mathbb{R}^{(n_1+1) \times (n_2+1)}, \\ &\quad P\mathbf{1}_{n_2+1} = (a_1, \dots, a_{n_1}, \sum_j b_j), \\ &\quad P^T\mathbf{1}_{n_1+1} = (b_1, \dots, b_{n_2}, \sum_i a_i)\}, \\ \Psi_C &:= \{(f, g) \in \mathbb{R}^{n_1+1} \times \mathbb{R}^{n_2+1}, \forall i, j, f_i + g_j \leq C_{ij}\}. \end{aligned}$$

Note that solving \mathbf{L}_C and recovering (primal or dual) optimizers can be done using, for instance, the **POT** library [FC17].

Motivations. Quantization can be fairly useful in practice from different perspectives.

- First, it allows us to “summarize” the information contained in a persistence diagram. Indeed, persistence diagrams in practice are likely to have a large number of points (e.g. few thousands). Each point accounts for a topological feature in the filtration process generating the diagram (see Section 2.1), but when the diagram is too big, this information is overloaded. Quantizing the diagram provides a persistence diagram (or measure) close to the input while being much simpler to interpret.
- Reducing the size of diagrams has huge benefits from a computational perspective. Computing the partial matching distance \mathbf{d}_p between two large diagrams is expensive. Given two diagrams μ, ν and a quantization Q , one can use $\text{OT}_p(Q(\mu), Q(\nu))$ as an estimation of $\text{OT}_p(\mu, \nu)$, where the error done is controlled by the quality of the quantization (using the triangle inequality).
- Finally, quantization can also be seen as a *vectorization* method (up to permutation). Indeed, an advantage of quantization in applications (in machine learning in particular) is that it turns a persistence diagram μ of arbitrary size into a vector in $\mathbb{R}^{3 \times k}$ (two dimensions for the point coordinates x_i , and one for the weight a_i), allowing for the use of many machine learning algorithms. Note that the k points in the (quantized) diagrams are not ordered (that is, permuting coordinates and weights of the vector in \mathbb{R}^3 should be considered as representing the same diagram); this can be handled by using *permutation invariant* layers such as PERSLAY (see Section 6.2).

EXAMPLE 7.11. Let $\mathcal{D}_k := \mathcal{M}_k \cap \mathcal{D}$ denote the space of persistence diagrams with exactly k points (supported on $\overline{\Omega}$). Let $\nu = \sum_{j=1}^n \delta_{y_j}$ be a persistence diagram. Assume for the sake of simplicity that $\text{pers}(y_1) > \text{pers}(y_2) > \dots > \text{pers}(y_n)$. Then,

$$\arg \min_{\mu \in \mathcal{D}_k} \mathbf{d}_2(\mu, \nu) = \sum_{i=1}^k \delta_{y_i},$$

that is, it is the diagram obtained from ν by keeping the k points of largest persistence. This quantization has the benefits of being straightforward to compute. Note that this technique, namely “only keeping points with highest persistence”, is routinely used in applications (e.g. [CCI⁺20, CO08]). It is reinterpreted here as a solution of a natural minimization problem.

EXAMPLE 7.12. Let us move temporarily out of the space of persistence diagrams. A well-known quantization problem—in the Wasserstein space $\mathcal{W}^2(\mathbb{R}^d)$, see Section 2.2—is the k -means problem. Indeed, consider a point cloud $\{y_1 \dots y_n\} \subset \mathbb{R}^d$ that is encoded by a probability measure $\nu = \frac{1}{n} \sum_{j=1}^n \delta_{y_j}$. Given an integer k , the k -means problem can be phrased as

$$\underset{x \in \mathbb{R}^{d \times k}, a \in \mathbb{R}_+^k, \sum_i a_i = 1}{\text{minimize}} \left\{ W_2^2(\mu(a, x), \nu) \right\},$$

where $\mu(a, x) := \sum_{i=1}^k a_i \delta_{x_i}$, and W_2 is the 2-Wasserstein distance.

REMARK 7.13. Example 7.11 essentially solves the problem of quantizing “in \mathcal{D} ”, and we will thus focus on proposing quantizations whose output is not a persistence diagram, in sense that mass a_i on a point x_i might not be integers. Considering this larger space of measures, which is made possible by connecting the diagram metrics \mathbf{d}_p to the optimal transport metrics OT_p , allows us to obtain significantly better quantized diagrams while preserving the benefits of quantization mentioned above (optimizing on a larger space gives better minima). Note also that, although we restrict our focus to persistence diagrams in this section, the techniques presented below would apply to more general measures (e.g. empirical estimators of the EPD, see Section 7.1).

7.2.2 Optimization problem and Algorithm

General problem formulation. Let us now fix a persistence diagram $\nu = \sum_{j=1}^n \delta_{y_j}$ and $k \in \mathbb{N}$. Let \mathcal{X} be a subset of \mathcal{M}_k . Our goal is to minimize

the objective

$$\begin{aligned} F : \mathbb{R}_+^k \times \mathbb{R}^{k \times 2} &\rightarrow \mathbb{R} \\ (a, X) &\mapsto \mathbf{L}_{C(X,Y)}(a, b), \end{aligned} \quad (7.18)$$

under the constraint $\mu(a, X) \in \mathcal{X}$, where $\mu(a, X) = \sum_{i=1}^k a_i \delta_{x_i}$. For a given X , we denote by $K_X := \{a, \mu(a, X) \in \mathcal{X}\}$, and symmetrically, for a given a , we define $K_a := \{X, \mu(a, X) \in \mathcal{X}\}$.

REMARK 7.14. *Quantizations problems, such as k -means (Example 7.12), are in general (NP) hard to solve. The objective functional (7.18) is in general not convex. We will thus focus in the following on providing algorithms that converges to a local minimum of F (just as the Lloyd algorithm returns a local optimum of the k -means problem).*

We propose to address this problem by *alternate minimization*: given an initialization $(a^{(0)}, X^{(0)})$, we build a sequence $(a^{(t)}, X^{(t)})_t$ by defining

$$\begin{aligned} a^{(t+1)} &= \min_{a \in K_X} F(a, X^{(t)}), \\ X^{(t+1)} &= \min_{X \in K_a} F(a^{(t+1)}, X). \end{aligned} \quad (7.19)$$

Obviously, such a sequence necessarily converges to a local minimum of F . The two minimization sub-problems will hopefully be simpler to deal with.

Mass update. Fix the location $X \in \mathbb{R}^{k \times 2}$, and consider $F_X : a \mapsto F(a, X)$. The dual formulation (7.17) yields

$$F_X(a) = \max_{f, g \in \Phi_C} \sum_{i=1}^{n_1} a_i f_i + f_\Delta \left(\sum_j b_j \right) + \sum_{j=1}^{n_2} b_j g_j + g_\Delta \left(\sum_i a_i \right) \quad (7.20)$$

Assume we have found optimal f^*, g^* . The *envelope theorem* states that a (sub-)gradient of F_X is given by

$$\nabla_a F_X = (f_i^* + g_\Delta^*)_{1 \leq i \leq k}. \quad (7.21)$$

Note that, as a supremum of convex (in fact, linear) functions, F_X is convex. Therefore, assuming K_X is closed and convex, F_X can be minimized using any projected gradient descent-based technique. For now, we only considered in our implementations the (naive) Euclidean gradient step and projection, that is for some learning rate $\lambda > 0$,

$$a \leftarrow \text{proj}_{K_X}(a - \lambda \nabla_a F_X), \quad (7.22)$$

where proj_{K_X} is the projection on K_X for the Euclidean, that is $\text{proj}_{K_X}(a') := \arg \min_{a'' \in K_X} \|a'' - a'\|_2^2$. More sophisticated techniques (e.g. accelerated gradient descent [Nes05]) are considered and will be incorporated in [Gudhi](#) if adapted. We also briefly mention the recent work of L.Chizat [Chi19] which considers optimization of discrete measures supported on a manifold (which is not $(\tilde{\Omega}, \rho)$)—it is likely that tools introduced in this work can be adapted to handle persistence diagrams and this will be investigated in future work.

Location update. Fix the weights $a \in \mathbb{R}^k$, and consider $F_a : X \mapsto F(a, X)$. The primal formulation (7.16) yields

$$\min_{X \in K_a} F_a(X) = \min_{X \in K_a} \min_{P \in \Pi(a, b)} \sum_{ij} P_{ij} C_{ij}. \quad (7.23)$$

Once again, we propose to solve this (non-convex) problem by alternating minimization between $X \in K_a$ and $P \in \Pi(a, b)$. The latter simply consists in solving the primal optimal transport problem. Now, given an optimal P^* , the former reads

$$\underset{X \in K_a}{\text{minimize}} \sum_{1 \leq i, j \leq k, n} P_{ij}^* d(x_i, y_j)^p + \sum_{i=1}^k P_{i\Delta}^* d(x_i, \mathbf{p}_{\partial\Omega}(x_i))^p, \quad (7.24)$$

where we recall that $\mathbf{p}_{\partial\Omega}(x_i)$ is the (orthogonal) projection of x_i onto the diagonal $\partial\Omega$.

In contrast to the update of masses, solving this minimization problem significantly depends on the nature of d, p , and K_a . We provide below some detailed examples.

The meta-algorithm summarizing masses and locations updates is described in Algorithm 5.

REMARK 7.15. *It is worth noting that Algorithm 5 can be straightforwardly adapted to **not** update either the mass or the locations (simplifying the loop).*

REMARK 7.16. *One can replace optimal primal P^* and dual (f^*, g^*) variables by their entropic regularized counterparts, see Section 2.2.4 for an introduction to entropic regularization of optimal transport, and Chapter 5 to see how it adapts to persistence measures. It can improve computational efficiency when k, n are getting large, but also the regularity of the algorithm, in particular with respect to the initialization. Indeed, without the*

Algorithm 5 Quantization of persistence diagrams.

Input: Integer k , initial masses $a^{(0)} \in \mathbb{R}^k$ and locations $X^{(0)} \in \overline{\Omega}^k$.
 Constraints K_a, K_X .

Output: A local minimum (a, X) of (7.18).
 (Init) $a \leftarrow a^{(0)}, X \leftarrow X^{(0)}$.

while Not converged (1) **do**

while Not converged (2) **do**

Compute (f^*, g^*) optimal in (7.17) using **POT**

Set $\nabla = (f_i^* + g_\Delta^*)_{i=1}^k$

$a \leftarrow a - \lambda \nabla$

$a \leftarrow \text{proj}_{K_a}(a)$

end while

while Not converged (3) **do**

Compute P^* optimal in (7.16) using **POT**

$X \leftarrow X^*$ optimal in (7.24).

end while

end while

return (a, X) .

entropic regularization, the optimal transport plan P^ , as a function of the current locations X , is not stable in sense that it is not unique (although being generically unique) and two close locations X', X'' can lead to two significantly different optimal transport plans $P^{*'}, P^{*''}$ and thus significant differences regarding the output of the algorithm. In contrast, the entropic regularized counterpart of P^* is unique and continuous with respect to X .*

7.2.3 Examples

Let us provide examples in three natural contexts in which one may want to perform quantization of persistence diagrams. In the following, we fix an integer k and an input measure $\nu(b, Y) = \sum_{j=1}^n b_j Y_j$, $b \in \mathbb{R}^n$, $Y \subset \Omega$.

Unconstrained case. The simplest case to consider is the unconstrained one, that is we put no constraint on either locations nor masses aside from having non-negative masses, i.e. belonging to the positive orthant \mathbb{R}_+^k . With previous notations, and with $q = p = 2$ —that is $d(x, y)^p = \|x - y\|_2^2$, it reads

- $a \geq 0$ entry-wise (formally, it yields $\forall X, K_X = \mathbb{R}_+^k$).

- $X \in \overline{\Omega}^k$. Let $X^{(0)}, X^{(1)} \in \mathbb{R}^k$ denote the first and second coordinates of X respectively. We must have $X^{(1)} \geq X^{(0)}$ (element-wise). However, one can observe that if there is a point $x = (x^{(0)}, x^{(1)})$ in X such that $x^{(1)} < x^{(0)}$ (the point is below the diagonal), then replacing x by $x' = (x^{(1)}, x^{(0)})$ in X will decrease the objective $F(a, X)$. Therefore, this constraint is not even required explicitly. Note that $K_a = \overline{\Omega}^k$ for all $a \in \mathbb{R}_+^k$

In this context, the mass update (7.22) requires only to project onto the non-negative orthant, which simply reads

$$\text{proj}_{\mathbb{R}_+^k}(x) = \max(x, 0),$$

where the maximum is meant term-wise. We thus focus on the location update. It simplifies to

$$\begin{aligned} \underset{x_1^{(0)} \dots x_k^{(0)}, x_1^{(1)} \dots x_k^{(1)} \in \mathbb{R}}{\text{minimize}} \quad & \sum_{ij} P_{ij}^* \left[\left(x_i^{(1)} - y_j^{(1)} \right)^2 + \left(x_i^{(0)} - y_j^{(0)} \right)^2 \right] \\ & + \frac{1}{2} \sum_i P_{i\Delta}^* \left(x^{(1)} - x^{(0)} \right)^2 + \sum_j P_{\Delta j}^* \left(y_j^{(1)} - y_j^{(0)} \right)^2. \end{aligned}$$

PROPOSITION 7.17. *Assume that $a > 0$ element-wise. For $1 \leq i \leq k$, let $\mathbf{P}_i := \sum_{j=1}^n P_{ij}^*$ and $\bar{y}_i := \frac{1}{\mathbf{P}_i} \sum_{j=1}^n P_{ij}^* y_j$. Also, let $\tau_i := \frac{\mathbf{P}_i}{a_i} \leq 1$. The optimality condition reads*

$$\begin{cases} x_i^{(0)} = x_i^{(1)} & \text{if } \mathbf{P}_i = 0, \\ x_i = \tau_i \bar{y}_i + (1 - \tau_i) \mathbf{p}_{\partial\Omega}(\bar{y}_i) & \text{otherwise.} \end{cases} \quad (7.25)$$

Here, \mathbf{P}_i represents the quantity of mass initially located on x_i that is off-diagonally transported, and \bar{y}_i represents a barycenter of the $(y_j)_j$ weighted by the proportion of mass P_{ij}^*/\mathbf{P}_i transported from x_i to y_j . The new position of x_i is eventually a barycenter between \bar{y}_i and its orthogonal projection onto $\partial\Omega$, where the weight τ_i represents the fraction of mass that is off-diagonally transported.

Proof. Let \mathcal{F} denote the function to minimize. Let $1 \leq i \leq k$, $\alpha \in \{0, 1\}$, and $\alpha' = 1 - \alpha$. The partial derivative of this function with respect to x_i^α is given by

$$\partial_{x_i^\alpha} \mathcal{F} = 2 \sum_{1 \leq j \leq n} P_{ij}^* (x_i^\alpha - y_j^\alpha) + P_{i\Delta}^* (x_i^\alpha - x_i^{\alpha'}).$$

Recall the relation $\mathbf{P}_i + P_{i\Delta} = a_i$ (marginal constraint on the transportation polytope). Equalizing this quantity to 0 gives

$$2\mathbf{P}_i x_i^\alpha - 2\mathbf{P}_i \bar{y}_i^\alpha + P_{i\Delta}(x_i^\alpha - x_i^{\alpha'}) = 0,$$

that is
$$a_i x_i^\alpha - \mathbf{P}_i \bar{y}_i^\alpha = P_{i,\Delta} \left(\frac{x_i^\alpha + x_i^{\alpha'}}{2} \right),$$

and same relation holds interchanging α and α' , so that we have

$$\begin{cases} x_i^\alpha + x_i^{\alpha'} = \bar{y}_i^\alpha + \bar{y}_i^{\alpha'} \\ a_i x_i^\alpha - \mathbf{P}_i \bar{y}_i^\alpha = a_i x_i^{\alpha'} - \mathbf{P}_i \bar{y}_i^{\alpha'} \end{cases} \Rightarrow \begin{cases} x_i^\alpha + x_i^{\alpha'} = \bar{y}_i^\alpha + \bar{y}_i^{\alpha'} \\ x_i^\alpha - x_i^{\alpha'} = \tau_i(\bar{y}_i^\alpha - \bar{y}_i^{\alpha'}) \end{cases}$$

which gives

$$\begin{cases} 2x_i^\alpha = (1 - \tau_i)(\bar{y}_i^\alpha + \bar{y}_i^{\alpha'}) + 2\tau_i \bar{y}_i^{\alpha'}, \\ 2x_i^{\alpha'} = (1 - \tau_i)(\bar{y}_i^{\alpha'} + \bar{y}_i^\alpha) + 2\tau_i \bar{y}_i^\alpha, \end{cases} \quad (7.26)$$

which proves the claim, using that $\mathbf{p}_{\partial\Omega}(\bar{y}_i) = \left(\frac{\bar{y}_i^{(0)} + \bar{y}_i^{(1)}}{2}, \frac{\bar{y}_i^{(0)} - \bar{y}_i^{(1)}}{2} \right)$. \square

Mass constrained. The unconstrained case presents the advantage of being fairly simple to solve. One might however want the quantized diagram to have more properties. For instance, one may ask that, although having k points, the quantized diagram should have the same total mass as the input diagram. Since the mass of the input diagram corresponds to its number of points (counted with multiplicity), and each point accounts for a topological event occurring in the filtration process, this constraint can be interpreted as asking for the quantized diagram to account for the same “total number of events” observed in the filtration process. Formally, the constraint on the masses a simply reads (for any location X)

$$\sum_{i=1}^k a_i = a^T \mathbf{1}_k = n,$$

that is, a must lie on a re-scaled version of the probability simplex of dimension k , $\{u \in \mathbb{R}^k, u^T \mathbf{1}_k = 1, u \geq 0\}$. The (Euclidean) projection onto the probability simplex can be found using, for instance, [WCP13, Algorithm 1]. Cuturi and Doucet [CD14, §4.2] propose to rely on a Bregman projected gradient descent. In that context, a gradient step reads

$$a \leftarrow n \frac{a \odot \exp(-\lambda \nabla_a F)}{\|a \odot \exp(-\lambda \nabla_a F_X)\|_1},$$

where \odot denotes element-wise multiplication, $\lambda > 0$ is a learning rate, and $\nabla_a F_X = (f_i^* + g_\Delta^*)_i$ for an optimal pair of dual potentials (f^*, g^*) .

Note that the locations remain unconstrained, and one can use Proposition 7.17 to update the locations.

Persistence constrained. We finally consider a more complicated variation where we ask the quantized diagram to share the same *total persistence* as the input diagram ν . Recall that for a parameter p , the total persistence of $\mu = \sum_{i=1}^k a_i \delta_{x_i}$ is defined as

$$\text{Pers}_p(\mu) = \sum_{i=1}^k a_i \text{pers}(x_i)^p = a^T \text{pers}(X),$$

where $\text{pers}(X) := (\text{pers}(x_i)^p)_i = (d(x_i, \mathbf{p}_{\partial\Omega}(x_i))^p)_i$. Informally, the total persistence can be understood as the total quantity of topological information contained in a diagram. Preserving the total persistence appears thus as a natural constraint.

Let $\rho = \text{Pers}_p(\nu)$ be the total persistence of the input diagram ν . For a given X , the constraint on a is given by $K_X = \{u \geq 0\} \cap \{u, u^T \text{pers}(X) = \rho\}$, which is the intersection of two convex sets: the non-negative orthant and an affine plan, it is thus convex. Furthermore, if X has at least one off-diagonal point, this intersection is not empty (assume $\text{pers}(x_i) > 0$, and take $a = (0 \dots 0, \rho/\text{pers}(x_i), 0 \dots 0)$). The (Euclidean) projection on this set can be found using a Linear Program solver.

Let us fix the mass a . The subtlety is that now, the locations $X = (x_i)_{i=1}^k$ are also constrained to satisfy $a^T \text{pers}(X) = \rho$. In particular, this constraint is not convex if we set $p = 2$, yielding $d(x, y)^p = \|x - y\|_2^2$. We would face a Quadratically Constrained Quadratic Program, with non-convex constraints. Such a problem cannot be handle by a standard optimizer solver and would require to do optimization on the manifold

$$\left\{ \frac{1}{2} \sum_{i=1}^n (x_i^{(1)} - x_i^{(0)})^2 = \rho \right\} = \{\mu \in \mathcal{M}_k, \text{OT}_2^2(\mu, \mathbf{0}) = \rho\},$$

which corresponds to a sphere of radius $\sqrt{\rho}$ centered on the empty diagram $\mathbf{0}$ on $(\mathcal{M}_k, \text{OT}_2)$. Optimization on manifold is a wide topic for which techniques exist [AMS09, Bou14]. Those have not been implemented at the time I am writing these lines and are left for future work.




Figure 7.3: Persistence constrained quantization with $k = 10$. (Left) Initial state of the quantization algorithm. Initialization on location is performed using a simple k -mean algorithm on the points of the input diagram ν (shaded red). Initialization of weights is with $a_i = \rho / \sum_i \text{pers}(x_i)$, so that $\text{Pers}_1(\mu(a_{\text{init}}, X_{\text{init}})) = \rho$. (Middle) Output of the algorithm after convergence with constrained persistence. (Right) Evolution of $\mathcal{E}_t := \text{OT}_1(\mu^{(t)}, \nu)$ through iterations t . Stopped when $|\mathcal{E}_{t+1}/\mathcal{E}_t - 1| \leq 0.001$.

In order to recover a convex problem, we instead consider in the following the case $p = q = 1$, that is $d(x, y)^p = \|x - y\|_1$. The constraints now read

$$\sum_{i=1}^k a_i(x_i^{(1)} - x_i^{(0)}) = \rho, \quad X^{(1)} \geq X^{(0)} \quad (7.27)$$

and are thus convex. Meanwhile, the objective functional reads, for an optimal transport plan P^* ,

$$\begin{aligned} F_a(X) &= \sum_{ij} P_{ij}^* (|x_i^{(0)} - y_j^{(0)}| + |x_i^{(1)} - y_j^{(1)}|) \\ &\quad + \sum_i P_{i\Delta}^*(x_i^{(1)} - x_i^{(0)}) + \sum_j P_{\Delta j}^*(y_j^{(1)} - y_j^{(0)}), \end{aligned}$$

which is convex. Minimizing F_a under the constraints (7.27) problem can be reformulated as a Linear Program, and thus finding an optimal X can be handled by a convex optimization solver such as [GLPK](#) interfaced with [cvxpy](#) [DB16], as done in our implementation. See Figure 7.3 for an illustration.

7.3 Shift-invariant distance between persistence diagrams

Motivations and preliminary remarks

Recall that a persistence diagram μ built on top of an object X comes from a *filtration* $f : X \rightarrow \mathbb{R}$, which is denoted by $\mu = \text{Dgm}(X, f)$ —see

Section 2.1 for details. In particular, the points $(x_i)_i$ of $\mu = \sum_{i=1}^n \delta_{x_i}$ are linked with some *critical values* of the map f . Roughly speaking², the coordinates $(b_i, d_i) \in \Omega$ of a point x_i in the diagram $\mu = \text{Dgm}(X, f)$ must correspond to local minima or maxima of f reached at b_i and d_i . As such, if one composes f with a (strictly) increasing map $\varphi : \mathbb{R} \rightarrow \mathbb{R}$, the resulting diagram $\text{Dgm}(X, \varphi \circ f)$ —that will also be denoted by $\varphi(\mu)$ —is simply described by $\sum_{i=1}^n \delta_{\varphi(x_i)}$, where $\varphi(x_i) := (\varphi(b_i), \varphi(d_i)) \in \Omega$, with $x_i = (b_i, d_i) \in \Omega$. Note that in particular $\text{Dgm}(X, \varphi \circ f)$ and $\text{Dgm}(X, f)$ have the same number of points n , and one has a natural correspondence between those points (although this might not be the correspondence given by an optimal partial matching between the two diagrams). To sum up, the two diagrams are the same up to a transformation of the ground space Ω .

EXAMPLE 7.18. *Let $X \subset \mathbb{R}^d$ be a (finite) point cloud centered at 0, and λX be the point cloud obtained from X by scaling it with a factor $\lambda > 0$. One can build the Čech diagrams (see Example A.10 for details) on top of X and λX , using distances in log-scale. This would respectively give two diagrams μ and $\nu = \tau(\mu)$, where $\tau = \log(\lambda)$. If the scaling factor λ is far from 1, μ and ν are two very different diagrams (in terms of the diagram metric \mathbf{d}_p). However, one could argue that describing topology of a point cloud should be scale-independent. It invites us to consider a distance between diagram up to translation, yielding (7.28) below. See Figure 7.4 for an illustration.*

A natural question is thus: given two diagrams μ, ν , is it possible to compute their distance up to such a transformation of the ground space? Note that allowing for *any* increasing transformation φ is too general: for instance if we let $\varphi_\lambda : t \mapsto \lambda t$ with $\lambda > 0$, one has $\inf_{\lambda, \lambda'} \mathbf{d}_p^p(\varphi_\lambda(\mu), \varphi_{\lambda'}(\nu)) = 0$ for any diagrams μ, ν (by taking λ, λ' arbitrarily small, so that the points of the diagrams collapse to $(0, 0) \in \partial\Omega$). Motivated by Example 7.18 and Remark 7.20 below, we restrict in this section to the study of φ being a translation, that is being parameterized by $\tau \in \mathbb{R}$ with $\varphi_\tau : t \mapsto t + \tau$. To alleviate notations, we write $\varphi_\tau(\mu) =: \tau(\mu)$ for $\tau \in \mathbb{R}$ and $\mu \in \mathcal{D}^p$. Observe that $\tau(\mu)$ is simply obtained from μ by shifting all the points by a translation (τ, τ) parallel to the diagonal. Note also that $d(\tau(x), \partial\Omega) = d(x, \partial\Omega)$ for all $x \in \Omega$ so that $\text{Pers}_p(\mu) = \text{Pers}_p(\tau(\mu))$ and thus \mathcal{D}^p is stable by $\mu \mapsto \tau(\mu)$. Furthermore, $\mu \sim \nu \Leftrightarrow \exists \tau, \mu = \tau(\nu)$ defines an equivalence relation on \mathcal{D}^p ,

²This can be made formal by observing that the persistence module of $(X, \varphi \circ f)$ indexed by $t \in \mathbb{R}$ is exactly the one of (X, f) indexed by $\varphi(t) \in \mathbb{R}$.




Figure 7.4: (Left) Two circles, the largest obtained by scaling the smallest by a factor $\lambda = 3$. (Middle) The standard Čech diagrams of both circles. (Right) The Čech diagrams in log-scale. They are obtained from each other by a translation (parallel to the diagonal) by a factor $\tau = \log(\lambda)$. As such, their distance up to translation is 0.

hence we let $\text{cl}(\mu) := \{\tau(\mu), \tau \in \mathbb{R}\}$, and we denote by \mathcal{D}^p/τ the quotient space $\{\text{cl}(\mu), \mu \in \mathcal{D}^p\}$.³

It invites us to define, for $\mu, \nu \in \mathcal{D}^p$,

$$\mathbf{d}_p^\tau(\text{cl}(\mu), \text{cl}(\nu)) := \inf_{\tau \in \mathbb{R}} \mathbf{d}_p(\mu, \tau(\nu)). \quad (7.28)$$

PROPOSITION 7.19. *For any $\tau \in \mathbb{R}, \mu, \nu \in \mathcal{D}^p$, we have $\mathbf{d}_p(\tau(\mu), \tau(\nu)) = \mathbf{d}_p(\mu, \nu)$, thus \mathbf{d}_p^τ is well-defined. Furthermore, it is a distance on \mathcal{D}^p/τ .*

Proof. The relation $\mathbf{d}_p(\tau(\mu), \tau(\nu)) = \mathbf{d}_p(\mu, \nu)$ follows from the fact that $d(\tau(x), \tau(y)) = d(x, y)$ and $d(x, \partial\Omega) = d(\tau(x), \partial\Omega)$ for all $x, y \in \Omega$ and that $\pi \mapsto \tau_\# \pi$ induces a bijection between $\text{Adm}(\mu, \nu)$ and $\text{Adm}(\tau(\mu), \tau(\nu))$; so eventually $C_p(\pi) = C_p(\tau_\# \pi)$ and thus $\mathbf{d}_p(\mu, \nu) = \mathbf{d}_p(\tau(\mu), \tau(\nu))$. Therefore, the right-and-side in (7.28) does not depend on the choice of $\mu \in \text{cl}(\mu)$ (resp. $\nu \in \text{cl}(\nu)$) and \mathbf{d}_p^τ is well-defined.

Let us prove that it is a distance. Non-negativity is clear. Symmetry follows from the previous point (and the symmetry of \mathbf{d}_p). Now, let $\mu, \nu \in \mathcal{D}^p$ and assume that $\mathbf{d}_p^\tau(\text{cl}(\mu), \text{cl}(\nu)) = 0$. If $\mu = \mathbf{0}$ (the empty diagram), this yields $\inf_\tau \text{Pers}_p(\tau(\nu)) = \text{Pers}_p(\nu) = 0$ thus $\nu = \mathbf{0} = \mu$. Now, let $(\tau_n)_n$ be such that $\mathbf{d}_p(\mu, \tau_n(\nu)) \rightarrow 0$. Assume that $(\tau_n)_n$ is not bounded. Up to extracting a subsequence, we can assume that $\tau_n \rightarrow +\infty$ (resp. $-\infty$). By the characterization of convergence in $(\mathcal{D}^p, \mathbf{d}_p)$ (Theorem 3.7), we know

³Formally, $\text{cl}(\mu)$ is the orbit of μ under the action of $(\mathbb{R}, +)$ on \mathcal{D} defined by $\tau \cdot \mu := \tau(\mu)$.




Figure 7.5: Rotation of the ground space used to make computations simpler.

that for each compact subset $K \subset \Omega$, $\tau_n(\nu)(K) \rightarrow \mu(K)$. On the other hand, $\tau_n(\nu)(K) \rightarrow 0$ since $\tau_n \rightarrow +\infty$ (resp. $-\infty$) and $\text{Pers}_p(\nu) < +\infty$. Thus, $\mu(K) = 0$ and $\mu = \mathbf{0}$, and $\nu = \mathbf{0}$ as well. Otherwise, it means that $(\tau_n)_n$ must be bounded, and thus up to extracting a subsequence, converges to some $\tau^* \in \mathbb{R}$. As $\tau_n(\nu) \rightarrow \tau^*(\nu)$, we obtain $\mathbf{d}_p(\mu, \tau^*(\nu)) = 0$, that is $\text{cl}(\mu) = \text{cl}(\nu)$. Finally, let $\mu, \nu, \lambda \in \mathcal{D}^p$. For all $\tau, \tau' \in \mathbb{R}$ we have

$$\begin{aligned} \mathbf{d}_p^\tau(\text{cl}(\mu), \text{cl}(\lambda)) &\leq \mathbf{d}_p(\mu, \tau'(\lambda)) \\ &\leq \mathbf{d}_p(\mu, \tau(\nu)) + \mathbf{d}(\tau(\nu), \tau'(\lambda)) \\ &= \mathbf{d}_p(\mu, \tau(\nu)) + \mathbf{d}(\nu, (\tau' - \tau)(\lambda)), \end{aligned}$$

taking the infimum in τ' , then in τ , gives the triangle inequality. \square

REMARK 7.20. *In symplectic topology, recent advances in the field were brought by considering persistence diagrams built using a Hamiltonian as a filtration. However, such filtrations are defined up to a constant term; the resulting persistence diagram is therefore only defined up to a translation parallel to the diagonal—that is an element of \mathcal{D}^p/τ , see [RSV18, Remark 34] or the paragraph **barcodes** in [BHS18, §1.1] for more details. The natural metric that arises in this context to compare these topological descriptors is then exactly $\mathbf{d}_2^{\tau,2}$.*

Computational aspects

REMARK 7.21. *For the ease of notations and computations, in the remainder this section we restrict to the case $p = 2$ and rotate the ground space Ω by $-\frac{\pi}{4}$. In particular, the distance between a point $x = (x^{(0)}, x^{(1)})$ and the diagonal $\partial\Omega$ is now simply given by $x^{(1)}$, see Figure 7.5.*




Figure 7.6: Illustration of the non-convexity of $\tau \mapsto \mathbf{d}_2^2(\mu, \tau(\nu))$ in a simple case. Here, $\mu = \delta_{x_1} + \delta_{x_2}$ (red points), and $\nu = \delta_y$ (blue point). Here, f has two (global) minima: $\tau = -1$ and $\tau = 1$. However, $f(0)$ is not a minimum. On the left, the graph of the map $\tau \mapsto f(\tau)$, which illustrates (a) non convexity of f , (b) presence of points of non-differentiability, (c) presence of plateaus.

REMARK 7.22. *In standard optimal transport (see Section 2.2) on the Euclidean space \mathbb{R}^d , the analog problem is trivial: one just has to center the measures considered. Namely, if μ, ν are two probability measures in \mathbb{R}^d . We have that⁴ $\inf_{\tau \in \mathbb{R}^d} W_2(\mu, \tau(\nu))$ is reached for $\tau = \mathbb{E}(\mu) - \mathbb{E}(\nu)$, where W_2 denotes the 2-Wasserstein distance on the Wasserstein space $\mathcal{W}^2(\mathbb{R}^d)$.*

However, including the diagonal, even if the set of translations considered here is only one-dimensional ($\tau \in \mathbb{R}$), the problem becomes harder as it is, in particular, not convex, see Figure 7.6.

Limitation of naive variational approaches. Intuitively, one could try to (locally) minimize the map $f : \tau \mapsto \mathbf{d}_2^2(\mu, \tau(\nu))$ by using either an alternate minimization or by doing a gradient descent as explained previously. Unfortunately, f is likely to have many plateaus (intervals on which it takes a constant value) or points of non-differentiability due to the peculiar role played by the diagonal, see for instance Figure 7.6. As a consequence, if initialized in such plateaus, both methods would get stuck and return a clearly unsatisfactory result.

REMARK 7.23. *We point out that a similar problem was already considered by Sheehy et al. in [SKC18]. Their method consists in adopting a combinatorial approach to provide a global minimizer of F . Their work differs from the one we present here in the following ways:*

- Sheehy et al. consider the bottleneck distance (case $p = \infty$, i.e. turning the sum into a max), while we consider the Wasserstein distance with

⁴With $\tau(\nu)$ the push-forward of ν by the map $x \mapsto x + \tau$.

$p = 2$ (although our approach works for any $p \geq 1$ using as ground metric $d(x, y) = \|x - y\|_p$).

- They did not focus on tractability of their algorithm: theoretical complexity is in $\mathcal{O}(n^{3.5})$ (where n is the cardinality of the diagrams in input) and they report a running time of 2 minutes for $n = 128$. Our implementation can handle much larger instances.
- Although developed in the context of persistence diagrams, our approach straightforwardly adapts to any discrete measure supported on Ω and, theoretically, even to measures with a non-discrete support, adopting the formalism developed in Chapter 3, allowing to compute shift-invariant distance between persistence surfaces (see Section 6.1 and [AEK⁺17]), expected persistence diagrams (see Section 7.1) or quantized diagrams (see Section 7.2) for instance.

Entropic optimal transport to get non-zero gradients

In order to mitigate the effect of these plateaus, we propose to use (once again) entropic optimal transport. This will also have benefits on differentiability and computational efficiency. With the notations above, let $X = (x_1 \dots x_n) \in R^{n \times 2}$ and $Y = (y_1 \dots y_m) \in R^{m \times 2}$ be the points appearing in μ and ν respectively. Define $a = (1 \dots 1, m)/(n + m) \in \mathbb{R}^{n+1}$ and $b = (1 \dots 1, n)/(n + m) \in \mathbb{R}^{m+1}$. Note that here, we renormalized a, b so that they belong to the probability simplex, allowing the use of standard (regularized) optimal transport results off-the-shelf.

We recall that we work with $p = 2$, and we thus we define $c(x, y) = \|x - y\|_2^2$ and $c(x, \partial\Omega) := \|x - \mathbf{p}_{\partial\Omega}(x)\|_2^2 = (x^{(1)})^2$, where $x = (x^{(0)}, x^{(1)}) \in \Omega$ (recall that we performed a rotation of the ground space Ω by $-\pi/4$ to simplify computations, see Figure 7.5). Consider $C(\tau) := c(x_i, y_j + (\tau, 0))$ the cost matrix where locations Y of ν get translated by τ , parallel to $\partial\Omega$. Note that $c(\partial\Omega, y) = c(\partial\Omega, y + (\tau, 0))$. Our goal becomes to minimize the function $f(\cdot, \gamma) : \mathbb{R} \rightarrow \mathbb{R}$ defined for a parameter $\gamma \geq 0$ by:

$$f(\tau, \gamma) = \min_{P \in \Pi(a, b)} E(P, \tau, \gamma) \quad (7.29)$$

with

$$E(P, \tau, \gamma) := \langle P, C(\tau) \rangle + \gamma \text{KL}(P|a \otimes b) \quad (7.30)$$

which corresponds to the primal formulation of entropic regularized optimal transport problem (see Section 2.2 and Chapter 5). Here, $\Pi(a, b)$ denotes the transportation polytope between a and b and $\text{KL}(P|a \otimes b) := \sum_{ij} P_{ij} \log \left(\frac{P_{ij}}{a_i b_j} \right)$ is the Kullback-Leibler divergence.




Figure 7.7: Illustration of the role of the regularization γ when $\mu = \nu = \delta_x$ is a diagram with a single point. When $\gamma = 0$, $f(\tau, 0) = \mathbf{d}_2^2(\mu, \tau(\nu))$ takes a constant value as long as $|\tau| > 2\|x - \mathbf{p}_{\partial\Omega}(x)\|_2^2$ —in which case it becomes cheaper to transport the point of both diagram to the diagonal. Increasing the value of γ helps to mitigate this phenomenon.

We denote by $P(\tau, \gamma)$ the minimizer of $P \mapsto E(P, \tau, \gamma)$, which is unique if $\gamma > 0$, and by convention $P(\tau, 0)$ denotes the minimizer of $P \mapsto E(P, \tau, 0)$ of largest entropy. It is well known [PC17, Prop. 4.1] that $P(\tau, \gamma) \rightarrow P(\tau, 0)$ as $\gamma \rightarrow 0$. We also denote by $P_{\partial\Omega}$ the “trivial” transport plan that maps every point to (its orthogonal projection onto) the diagonal $\partial\Omega$, formally defined as

$$P_{\partial\Omega} := \begin{pmatrix} & & & \frac{1}{n+m} \\ & & & \vdots \\ & 0 & & \frac{1}{n+m} \\ \frac{1}{n+m} & \dots & \frac{1}{n+m} & 0 \end{pmatrix} \in \mathbb{R}^{(n+1) \times (m+1)}.$$

By definition, we know that for all τ, γ ,

$$\begin{aligned} f(\tau, \gamma) &\leq E(P_{\partial\Omega}, \tau, \gamma) = \frac{1}{n+m} (\text{Pers}(\mu) + \text{Pers}(\nu)) + \gamma \text{KL}(P_{\partial\Omega} | a \otimes b) \\ &=: E_{\partial\Omega}(\gamma), \end{aligned}$$

where $\text{Pers}(\mu) = \sum_{i=1}^n \|x_i - \mathbf{p}_{\partial\Omega}(x_i)\|_2^2$ is the *total persistence* of μ ($\text{Pers}(\nu)$ is defined likewise). Importantly, the latter quantity does not depend on τ , and thus $f(\cdot, \gamma)$ is bounded for all $\gamma \geq 0$.

Some intuition. Why could it help to add this entropic regularization? Aside from its general strengths (faster computations, access to gradients,

etc.), in our context, the key idea is that the (entropic regularized) optimal transport plan $P(\tau, \gamma)$ is *dense*, that is $P(\tau, \gamma)_{ij} > 0$ for all i, j . To be more precise, there exist two diagonal matrices u, v (unique up to a scaling factor) such that

$$P(\tau, \gamma) = u \cdot \exp\left(-\frac{C(\tau)}{\gamma}\right) \cdot v.$$

Therefore, in contrast to the non-regularized optimal transport plan (which is, in this case, a permutation matrix (Proposition 2.20 and Proposition 3.5)), the smoothed transport plan takes account of interactions between *all* pairs of points x, y even if those are far away. In particular, the situation where everything is transported onto the diagonal $\partial\Omega$ cannot occur when adding this entropic regularization⁵. See Figure 7.7 for an illustration.

We now state asymptotic properties (in both γ and τ) of f .

PROPOSITION 7.24 (Asymptotic behavior). *One has*

$$\lim_{\gamma \rightarrow 0} f(\tau, \gamma) = \min_{P \in \Pi(a, b)} \langle P, C(\tau) \rangle =: f(\tau, 0). \quad (7.31)$$

and

$$\lim_{\gamma \rightarrow \infty} f(\tau, \gamma) = \langle a \otimes b, C(\tau) \rangle =: f(\tau, \infty). \quad (7.32)$$

In particular, $f(\tau, \infty)$ admits a unique minimizer

$$\tau(\infty) := \frac{1}{m} \sum_j y_j^0 - \frac{1}{n} \sum_i x_i^0. \quad (7.33)$$

Similarly, one has

$$\lim_{|\tau| \rightarrow \infty} f(\tau, \gamma) = E_{\partial\Omega}(\gamma). \quad (7.34)$$

Proof. Equations (7.31) and (7.32) are now standard results in optimal transport theory, see [Gen19, Theorem 10] for instance. Note that

$$\begin{aligned} \langle a \otimes b, C(\tau) \rangle &= \frac{1}{(n+m)^2} \sum_{1 \leq i, j \leq n, m} [(\tau + y_j^{(0)} - x_i^{(0)})^2 + (y_j^{(1)} - x_i^{(1)})^2] \\ &\quad + \frac{n}{(n+m)^2} \text{Pers}(\nu) + \frac{m}{(n+m)^2} \text{Pers}(\mu), \end{aligned}$$

which give the value of $\tau(\infty)$.

⁵Note that however, in practice, if $C(\tau)/\gamma$ has some large entries at indices i, j , it numerically yields $\exp(-C(\tau)_{ij}/\gamma) \simeq 0$, which would impose $P_{ij} = 0$.

For the asymptotic in τ , let $P(\tau, \gamma)$ denote the (unique) minimizer of f . Assume that for some $1 \leq i, j \leq n, m$, we do not have $P_{ij}(\tau, \gamma)C_{ij}(\tau)$ bounded as $\tau \rightarrow \infty$. As the Kullback-Leibler divergence is always non-negative, we have $f(\tau, \gamma) \geq \sum_{ij} C_{ij}(\tau)P_{ij}(\tau, \gamma) \rightarrow \infty$. This contradicts the fact that f must remain bounded. Therefore, as $C_{ij}(\tau) \rightarrow \infty$, we necessarily have $P_{ij}(\tau, \gamma) \rightarrow 0$ and thus $P(\tau, \gamma) \rightarrow P_{\partial\Omega}$ (due to marginal constraints) and (as E is continuous) $f(\tau, \gamma) \rightarrow E_{\partial\Omega}(\gamma)$. \square

Intuitively, the larger the regularization γ , the simpler it is to minimize $f(\cdot, \gamma)$, to the extreme case $\gamma = \infty$ where $f(\cdot, \infty)$ is even convex with a minimizer $\tau(\infty)$ that can be found in closed-form. Of course, a (global) minimizer $\tau(\gamma)$ of $f(\cdot, \gamma)$ has no reason to be a good approximation of some $\tau(0)$ (a minimizer of our initial objective, $f(\tau, 0)$). There is a natural trade-off between a large γ (making the optimization problem simple) and a small one, that would give satisfactory minimizers. We first make this last point more precise.

PROPOSITION 7.25 (Convergence of minima and minimizers). *We have convergence of minimum values:*

$$\lim_{\gamma \rightarrow 0} \min_{\tau \in \mathbb{R}} f(\tau, \gamma) = \min_{\tau \in \mathbb{R}} f(\tau, 0). \quad (7.35)$$

Now, let $\tau(\gamma)$ be in $\arg \min f(\tau, \gamma)$ for $\gamma > 0$. One has

$$\begin{cases} \lim_{\gamma \rightarrow 0} \tau(\gamma) \in \arg \min_{\tau} f(\tau, 0) \\ f(\tau(\gamma), \gamma) \rightarrow \min_{\tau} f(\tau, 0) \end{cases} \quad (7.36)$$

In order to prove this proposition, we introduce the following lemma.

LEMMA 7.26 (Variations). *One has, for any (τ, γ) and (τ', γ') :*

$$|f(\tau, \gamma) - f(\tau, \gamma')| \leq 2|\gamma - \gamma'| \log(n + m). \quad (7.37)$$

In particular, the map $f(\cdot, \gamma)$ uniformly converges to $f(\cdot, 0)$ as $\gamma \rightarrow 0$.

Similarly,

$$|f(\tau', \gamma) - f(\tau, \gamma)| \leq \frac{1}{n + m} \|C(\tau) - C(\tau')\|_1 \quad (7.38)$$

Proof. Recall that $P(\tau, \gamma)$ denotes the unique minimizer of $P \mapsto E(P, \tau, \gamma)$ (or, in the particular case $\gamma = 0$, the one with largest entropy).

By definition,

$$\begin{aligned} f(\tau, \gamma) &= \langle P(\tau, \gamma), C(\tau) \rangle + \gamma \text{KL}(P(\tau, \gamma)|a \otimes b) \\ &\leq \langle P(\tau, \gamma'), C(\tau) \rangle + \gamma \text{KL}(P(\tau, \gamma')|a \otimes b) \\ &= f(\tau, \gamma') + (\gamma - \gamma') \text{KL}(P(\tau, \gamma')|a \otimes b). \end{aligned}$$

Now, observe that $\text{KL}(P|a \otimes b) = \sum_{ij} P_{ij} \log(P_{ij}) - \sum_{ij} P_{ij} \log(a_i b_j)$. The first term is non-positive. Furthermore,

$$\begin{aligned} - \sum_{ij} P_{ij} \log(a_i b_j) &= \sum_{1 \leq i, j \leq n, m} P_{ij} \log((n+m)^2) \\ &\quad + \sum_{1 \leq i \leq n} P_{i, m+1} \log((n+m)^2/n) \\ &\quad + \sum_{1 \leq j \leq m} P_{n+1, j} \log((n+m)^2/m) \\ &\quad + P_{n+1, m+1} \log((n+m)^2/nm) \\ &= 2 \log(n+m) - (\log(n) \sum_{1 \leq i \leq n} P_{i, m+1} \\ &\quad + \log(m) \sum_{1 \leq j \leq m} P_{n+1, j} + \log(nm) P_{n+1, m+1}) \\ &= 2 \log(n+m) - \frac{n}{n+m} \log(n) - \frac{m}{n+m} \log(m) \\ &\leq 2 \log(n+m). \end{aligned}$$

It proves the first inequality.

The second inequality is obtained in a similar way. By definition,

$$\begin{aligned} f(\tau, \gamma) &\leq \langle P(\tau', \gamma), C(\tau) \rangle + \gamma \text{KL}(P(\tau', \gamma)|a \otimes b) \\ &= f(\tau', \gamma) + \langle P(\tau', \gamma), C(\tau) - C(\tau') \rangle. \end{aligned}$$

Note that for $i = n+1$ or $j = m+1$, C does not depend on τ . Using the Hölder inequality on the indices $1 \leq i, j \leq n, m$, we can write

$$\langle P(\tau', \gamma), C(\tau) - C(\tau') \rangle \leq \|P\|_q \|C(\tau) - C(\tau')\|_{q'}$$

where $1/q + 1/q' = 1$. Taking $q = \infty$, and observing that (on the aforementioned indices) $P(\cdot) \leq \frac{1}{n+m}$ (due to the marginal constraints) gives the result. \square

Proof of Proposition 7.25. Given uniform convergence of the maps (Lemma 7.26), convergences of minima and minimizers follow from standard result in Γ -convergence literature (see [DM12]). More precisely, (7.35) is given by [DM12, Theorem 7.4], and (7.36) is given by [DM12, Corollary 7.20]. \square

Proposition 7.25 supports that it is reasonable to minimize $f(\cdot, \gamma)$ for some small $\gamma > 0$. Actually, it even suggests that one can start by minimizing $f(\cdot, \gamma)$ with some large γ , obtaining a minimizer $\tau(\gamma)$, and progressively decreasing γ . These techniques cannot guarantee to reach a global optimum of $f(\cdot, 0)$, but significantly improve on doing a naive alternate minimization on $f(\cdot, 0)$ directly. We give in the following the formula to optimize τ doing either an alternate minimization or a gradient step. Recall that computing $P(\tau, \gamma)$ can be done using the Sinkhorn algorithm—see Section 2.2—and is handle in practice by the library `POT` [FC17] (for instance).

PROPOSITION 7.27 (Optimization steps). *Let $\tau \in \mathbb{R}$ and $\gamma > 0$. Let $P \in \Pi(a, b)$ be a transport plan and assume that $\sum_{1 \leq i, j \leq n, m} P_{ij} > 0$ (i.e. there is some off-diagonal mass that is transported). Then, alternate minimization can be performed by using the following formula:*

$$\arg \min_{\tau} E(P, \tau, \gamma) = \frac{\sum_{1 \leq i, j \leq n, m} P_{ij} (x_i^{(0)} - y_j^{(0)})}{\sum_{1 \leq i, j \leq n, m} P_{ij}}. \quad (7.39)$$

Now, let $P(\tau, \gamma)$ be the unique minimizer of $P \mapsto E(P, \tau, \gamma)$. Then gradient descent can be performed by using the following formula

$$\frac{\partial f(\tau, \gamma)}{\partial \tau} = 2 \sum_{1 \leq i, j \leq n, m} P(\tau, \gamma)_{ij} (\tau + y_j^{(0)} - x_i^{(0)}). \quad (7.40)$$

Proof. Equation (7.39) follows from straightforward computations. Indeed,

$$\begin{aligned} E(P, \tau, \gamma) &= \sum_{ij} P_{1 \leq ij \leq n, m} [(\tau + y_j^{(0)} - x_i^{(0)})^2 + (y_j^{(1)} - x_i^{(1)})^2] \\ &\quad + \sum_{i=1}^n P_{i, m+1} \|x_i - \mathbf{p}_{\partial\Omega}(x_i)\|_2^2 + \sum_{j=1}^m P_{n+1, j} \|y_j - \mathbf{p}_{\partial\Omega}(y_j)\|_2^2 \\ &\quad + \gamma \text{KL}(P|a \otimes b) \end{aligned}$$

which is a quadratic (in particular, convex) map in τ (as long as $\exists(1 \leq i, j \leq n, m) | P_{ij} > 0$) which is minimized for

$$\sum_{ij} P_{1 \leq ij \leq n, m} (\tau + y_j^{(0)} - x_i^{(0)}) = 0,$$

which proves the formula.

Equation (7.40) is obtained using the envelope theorem which here simply reads

$$\begin{aligned} \frac{\partial}{\partial \tau} f(\tau, \gamma) &= \sum_{ij} P(\tau, \gamma)_{ij} \frac{d}{d\tau} C(\tau)_{ij} \\ &= 2 \sum_{1 \leq i, j \leq n, m} P(\tau, \gamma)_{ij} (\tau + y_j^{(0)} - x_i^{(0)}), \end{aligned}$$

concluding the proof. \square

This approach is eventually summarized in Algorithm 6. We propose to initialize on $\tau(\infty)$ as defined in (7.33). We initialize γ at 1, but other values might be satisfactory too. We run our experiments with a preset series of decreasing values of γ for illustration purpose but other heuristics such as $\gamma \leftarrow \gamma/2$ would make sense. Note also that we have access to $\frac{\partial f}{\partial \gamma}$, which is simply given by $\text{KL}(P(\tau, \gamma) | a \otimes b)$ (once again, this follows from the envelope theorem).

Algorithm 6 Shift-invariant distance between persistence diagram.

Input: Initial diagrams locations X (fixed) and Y (optimized). Update rule for γ . If gradient descent, learning rate $\lambda > 0$.

Output: A local minimum τ^* of $f(\tau, \gamma)$.

(Init) $\tau = \tau(\infty)$, $\gamma = 1$.

while Not converged (1) **do**

while Not converged (2) **do**

 Compute $P(\tau_t, \gamma_{t'})$ optimal using **POT**.

if Gradient descent **then**

 Set $\nabla = \frac{\partial f}{\partial \tau}$ using (7.40).

$\tau_{t+1} \leftarrow \tau_t - \lambda \nabla$.

end if

if Alternate minimization **then**

 Set $\tau_{t+1} \leftarrow \arg \min_{\tau} E(P(\tau_t, \gamma_{t'}), \tau, \gamma_{t'})$ using (7.39).

end if

end while

 Update γ

end while

return (a, X) .




Figure 7.8: Illustration of the shift-invariant distance algorithm on a simple example. (Left) The two diagrams we want to “align”, at $\tau = 0$. Intuitively, one should shift the red diagram by $\tau \simeq -0.6$; however the presence of a red cluster close to some of the blue points makes $\tau \simeq 0$ a local optimum of $\tau \mapsto \mathbf{d}_2^2(\mu, \tau(\nu)) =: f(\tau, 0)$. (Right) The graphs of $\tau \mapsto f(\tau, \gamma)$ for different regularization parameters γ . In red dots, the values $(\tau(\gamma_t), f(\tau(\gamma_t), \gamma_t))$ for γ_t decreasing from 1.0 to 0.01. In this (lucky) configuration, our approach does converge (close) to the global optimum, while directly working with $\gamma = 0$ (or even a small γ) would get stuck in the local optimum around $\tau = 0$.

CONCLUSION

The similarities between the metrics used to compare persistence diagrams in topological data analysis and those used to compare probability measures in optimal transport were known for long. However, an explicit connection between these two formalisms was yet to come. In this thesis, we established such a connection. It turned out to be very prolific, especially when targeting geometric and statistical applications but also in the conception of new algorithms based on optimal transport techniques that could be of great use when dealing with persistence diagrams.

The contributions of this work were presented in two different parts. Part I is dedicated to the theoretical results. Chapter 3 is the engine of the machinery we developed. We show how persistence diagram metrics can be reformulated as optimal partial transport problems. Of major interest, our formalism handles any kind of non-negative measures, not only those with discrete support such as persistence diagrams. We study the properties of the resulting space in detail, proving some powerful results such as a characterization of convergence in the space. Chapter 4 studies Fréchet means in this space from a theoretical perspective. We have established a strong connection between Fréchet means of persistence diagrams and their counterpart in optimal transport theory. It allows us to prove their existence as well as a consistency result in a great generality.

Part II is dedicated to applications. Following Chapter 4, Chapter 5 is dedicated to the estimation of Fréchet means of persistence diagrams in practice. By switching from a Lagrangian approach to a Eulerian one, we retrieved a convex problem, a major improvement over previous approaches. Furthermore, building on entropic optimal transport, we obtain a GPU-friendly parallelizable implementation, especially efficient on large sets of large diagrams. Chapter 6 is dedicated to linear representations of persistence diagrams, for which we provide an exhaustive criterion to prove their continuity. We then propose a neural network architecture, named PERSLAY, that will *learn* a task-driven linear representation. Our

publicly available implementation is made to encompass most preexisting vectorization techniques used when incorporating persistence diagrams in machine learning pipelines. Finally, Chapter 7 provides some complementary examples that showcase the benefits of our work. We study persistence diagrams in a random setting for which measures with continuous support naturally appear. We prove stability properties of persistence diagrams in this context, a result that only makes sense after extending diagram metrics to measures with continuous support. This is done thanks to our optimal transport based approach. Similarly, we propose an algorithm to perform quantization on persistence diagrams, and a method to estimate a shift-invariant distance between them.

We believe this work opens the way for faithful use of optimal transport tools in combination with topological descriptors such as persistence diagrams, both in theory and in practice. Of course, many tracks remain to be explored. Pretty much any aspect of the optimal transport literature may lead to an interesting counterpart in topological data analysis. We think for instance of gradient flows [AGS08], convergence rates based on geometric properties of the space [GPRS19] and using different regularization terms [BSR18], to name a few.

BIBLIOGRAPHY

- [AC11] Martial Agueh and Guillaume Carlier. Barycenters in the wasserstein space. *SIAM Journal on Mathematical Analysis*, 43(2):904–924, 2011. *Cited (3) on pages 50, 92, and 93*
- [AEK⁺17] Henry Adams, Tegan Emerson, Michael Kirby, Rachel Neville, Chris Peterson, Patrick Shipman, Sofya Chepushtanova, Eric Hanson, Francis Motta, and Lori Ziegelmeier. Persistence images: a stable vector representation of persistent homology. *Journal of Machine Learning Research*, 18(8):1–35, 2017. *Cited (12) on pages 16, 27, 45, 127, 128, 134, 136, 140, 141, 142, 152, and 172*
- [AGS08] Luigi Ambrosio, Nicola Gigli, and Giuseppe Savaré. *Gradient flows: in metric spaces and in the space of probability measures*. Springer Science & Business Media, 2008. *Cited on page 182*
- [AKP17] Stephanie Alexander, Vitali Kapovitch, and Anton Petrunin. Alexandrov geometry. *Book in preparation*, 2017. *Cited on page 223*
- [AMS09] P-A Absil, Robert Mahony, and Rodolphe Sepulchre. *Optimization algorithms on matrix manifolds*. Princeton University Press, 2009. *Cited on page 166*
- [AWR17] Jason Altschuler, Jonathan Weed, and Philippe Rigollet. Near-linear time approximation algorithms for optimal transport via sinkhorn iteration. In *Advances in Neural Information Processing Systems*, pages 1961–1971, 2017. *Cited (3) on pages 119, 120, and 123*

- [BBB⁺01] Dmitri Burago, Iu D Burago, Yuri Burago, Sergei A Ivanov, and Sergei Ivanov. *A course in metric geometry*, volume 33. American Mathematical Soc., 2001. *Cited on page 223*
- [BD17] Peter Bubenik and Paweł Dłotko. A persistence landscapes toolbox for topological statistics. *Journal of Symbolic Computation*, 78:91–114, 2017. *Cited on page 136*
- [BE19] Peter Bubenik and Alex Elchesen. Universality of persistence diagrams and the bottleneck and wasserstein distances. *arXiv preprint arXiv:1912.02563*, 2019. *Cited on page 79*
- [BGMP14] Andrew J Blumberg, Itamar Gal, Michael A Mandell, and Matthew Pancia. Robust statistics, hypothesis testing, and confidence intervals for persistent homology on metric measure spaces. *Foundations of Computational Mathematics*, 14(4):745–789, 2014. *Cited on page 75*
- [BHS18] Lev Buhovsky, Vincent Humiliere, and Sobhan Seyfaddini. The action spectrum and c^0 symplectic topology. *arXiv preprint arXiv:1808.09790*, 2018. *Cited on page 170*
- [BLPY19] Greg Bell, Austin Lawson, C Neil Pritchard, and Dan Yasaki. The space of persistence diagrams fails to have yu 's property a. *arXiv preprint arXiv:1902.02288*, 2019. *Cited on page 27*
- [BM19] Christophe AN Biscio and Jesper Møller. The accumulated persistence function, a new useful functional summary statistic for topological data analysis, with a view to brain artery trees and spatial point process applications. *Journal of Computational and Graphical Statistics*, (just-accepted):1–20, 2019. *Cited on page 45*
- [BMEWL11] Thierry Bertin-Mahieux, Daniel P.W. Ellis, Brian Whitman, and Paul Lamere. The million song dataset. In *Proceedings of the 12th International Conference on Music Information Retrieval (ISMIR 2011)*, 2011. *Cited (2) on pages 9 and 20*
- [Boc33] Salomon Bochner. Integration von funktionen, deren werte die elemente eines vektorraumes sind. *Fundamenta Mathematicae*, 20(1):262–176, 1933. *Cited on page 153*
- [Bou14] Nicolas Boumal. *Optimization and estimation on manifolds*. PhD thesis, Catholic University of Louvain, Louvain-la-Neuve, Belgium, 2014. *Cited on page 166*

- [BSR18] Mathieu Blondel, Vivien Seguy, and Antoine Rolet. Smooth and sparse optimal transport. In *International Conference on Artificial Intelligence and Statistics*, pages 880–889, 2018.
Cited on page 182
- [BT03] Amir Beck and Marc Teboulle. Mirror descent and non-linear projected subgradient methods for convex optimization. *Operations Research Letters*, 31(3):167–175, 2003.
Cited on page 126
- [Bub15] Peter Bubenik. Statistical topological data analysis using persistence landscapes. *The Journal of Machine Learning Research*, 16(1):77–102, 2015.
Cited (5) on pages 16, 27, 45, 136, and 140
- [BV18] Peter Bubenik and Tane Vergili. Topological spaces of persistence modules and their properties. *Journal of Applied and Computational Topology*, pages 1–37, 2018.
Cited (3) on pages 27, 40, and 74
- [BW20] Peter Bubenik and Alexander Wagner. Embeddings of persistence diagrams into hilbert spaces. *Journal of Applied and Computational Topology*, 4(3):339–351, 2020.
Cited (2) on pages 16 and 27
- [CB15] William Crawley-Boevey. Decomposition of pointwise finite-dimensional persistence modules. *Journal of Algebra and its Applications*, 14(05):1550066, 2015.
Cited on page 213
- [CB19] Mathieu Carrière and Ulrich Bauer. On the metric distortion of embedding persistence diagrams into separable hilbert spaces. In *35th International Symposium on Computational Geometry (SoCG 2019)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2019.
Cited (2) on pages 16 and 27
- [CCI+20] Mathieu Carrière, Frédéric Chazal, Yuichi Ike, Théo Lacombe, Martin Royer, and Yuhei Umeda. Perslay: A neural network layer for persistence diagrams and new graph topological signatures. In *International Conference on Artificial Intelligence and Statistics*, pages 2786–2796. PMLR, 2020.
Cited (7) on pages 19, 29, 133, 138, 139, 141, and 160
- [CCO17] Mathieu Carrière, Marco Cuturi, and Steve Oudot. Sliced wasserstein kernel for persistence diagrams. In

- 34th International Conference on Machine Learning*, 2017.
Cited (8) on pages 16, 27, 44, 69, 136, 140, 141, and 142
- [CCSG⁺09a] Frédéric Chazal, David Cohen-Steiner, Marc Glisse, Leonidas J Guibas, and Steve Y Oudot. Proximity of persistence modules and their diagrams. In *Proceedings of the twenty-fifth annual symposium on Computational geometry*, pages 237–246. ACM, 2009. Cited (2) on pages 35 and 214
- [CCSG⁺09b] Frédéric Chazal, David Cohen-Steiner, Leonidas J Guibas, Facundo Mémoli, and Steve Y Oudot. Gromov-hausdorff stable signatures for shapes using persistence. In *Computer Graphics Forum*, volume 28, pages 1393–1403. Wiley Online Library, 2009. Cited on page 219
- [CD14] Marco Cuturi and Arnaud Doucet. Fast computation of wasserstein barycenters. In *International Conference on Machine Learning*, pages 685–693, 2014. Cited (8) on pages 50, 112, 114, 115, 117, 125, 126, and 165
- [CDSGO16] Frédéric Chazal, Vin De Silva, Marc Glisse, and Steve Oudot. *The structure and stability of persistence modules*. Springer, 2016. Cited (4) on pages 32, 213, 214, and 218
- [CDSO14] Frédéric Chazal, Vin De Silva, and Steve Oudot. Persistence stability for geometric complexes. *Geometriae Dedicata*, 173(1):193–214, 2014. Cited on page 142
- [CE10] G. Carlier and I. Ekeland. Matching for teams. *Economic Theory*, 42(2):397–418, 2010. Cited on page 92
- [CFG⁺15] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015. Cited (2) on pages 9 and 20
- [CFL⁺14] Frédéric Chazal, Brittany Terese Fasy, Fabrizio Lecci, Alessandro Rinaldo, and Larry Wasserman. Stochastic convergence of persistence landscapes and silhouettes. In *Proceedings of the thirtieth annual symposium on Computational geometry*, page 474. ACM, 2014. Cited (4) on pages 134, 136, 140, and 152

- [CFL⁺15] Frédéric Chazal, Brittany Fasy, Fabrizio Lecci, Bertrand Michel, Alessandro Rinaldo, and Larry Wasserman. Sub-sampling methods for persistent homology. In *International Conference on Machine Learning*, pages 2143–2151, 2015. *Cited on page 151*
- [CFP01] Francesca Cagliari, Massimo Ferri, and Paola Pozzi. Size functions from a categorical viewpoint. *Acta Applicandae Mathematica*, 67(3):225–235, 2001. *Cited on page 22*
- [CHF12] Matthew A Cleveland, John M Hickey, and Selma Forni. A common dataset for genomic analysis of livestock populations. *G3: Genes, Genomes, Genetics*, 2(4):429–435, 2012. *Cited (2) on pages 9 and 20*
- [Chi17] Lenaïc Chizat. *Transport Optimal de mesures positives: modèles, méthodes numériques, applications*. PhD thesis, Université Paris-Dauphine, 11 2017. *Cited on page 50*
- [Chi19] Lenaïc Chizat. Sparse optimization on measures with over-parameterized gradient descent. *arXiv preprint arXiv:1907.10300*, 2019. *Cited on page 162*
- [CO08] Frédéric Chazal and Steve Yann Oudot. Towards persistence-based reconstruction in euclidean spaces. In *Proceedings of the twenty-fourth annual symposium on Computational geometry*, pages 232–241. ACM, 2008. *Cited on page 160*
- [COO15a] Guillaume Carlier, Adam Oberman, and Edouard Oudet. Numerical methods for matching for teams and wasserstein barycenters. *ESAIM: Mathematical Modelling and Numerical Analysis*, 49(6):1621–1642, 2015. *Cited (3) on pages 92, 104, and 114*
- [COO15b] Mathieu Carrière, Steve Y. Oudot, and Maks Ovsjanikov. Stable topological signatures for points on 3d shapes. *Computer Graphics Forum*, 34(5):1–12, 2015. *Cited on page 128*
- [CP16] Marco Cuturi and Gabriel Peyré. A smoothed dual approach for variational wasserstein problems. *SIAM Journal on Imaging Sciences*, 9(1):320–343, 2016. *Cited on page 50*

- [CPSV18] L ena ic Chizat, Gabriel Peyr e, Bernhard Schmitzer, and Fran ois-Xavier Vialard. Unbalanced optimal transport: Dynamic and kantorovich formulations. *Journal of Functional Analysis*, 274(11):3090–3123, 2018. *Cited on page 50*
- [CSEH07] David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Stability of persistence diagrams. *Discrete & Computational Geometry*, 37(1):103–120, 2007. *Cited on page 35*
- [CSEHM10] David Cohen-Steiner, Herbert Edelsbrunner, John Harer, and Yuriy Mileyko. Lipschitz functions have 1 p-stable persistence. *Foundations of computational mathematics*, 10(2):127–139, 2010. *Cited (3) on pages 37, 149, and 150*
- [Cut13] Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *Advances in Neural Information Processing Systems*, pages 2292–2300, 2013. *Cited (2) on pages 54 and 124*
- [CWRW15] Yen-Chi Chen, Daren Wang, Alessandro Rinaldo, and Larry Wasserman. Statistical analysis of persistence intensity functions. *arXiv preprint arXiv:1510.02502*, 2015. *Cited on page 136*
- [DB16] Steven Diamond and Stephen Boyd. CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*, 17(83):1–5, 2016. *Cited on page 167*
- [DC19] Vincent Divol and Fr ed eric Chazal. The density of expected persistence diagrams and its kernel based estimation. *JoCG*, 10(2):127–153, 2019. *Cited (2) on pages 43 and 148*
- [DGK18] P Dvurechensky, A Gasnikov, and A Kroshnin. Computational optimal transport: Complexity by accelerated gradient descent is better than by sinkhorn’s algorithm. In *35th International Conference on Machine Learning, ICML 2018*, pages 2196–2220, 2018. *Cited on page 119*
- [DL19] Vincent Divol and Th eo Lacombe. Understanding the topology and the geometry of the persistence diagram space via optimal partial transport. *arXiv preprint arXiv:1901.03048*, 2019. *Cited (5) on pages 18, 19, 28, 29, and 148*

- [DM12] Gianni Dal Maso. *An introduction to Γ -convergence*, volume 8. Springer Science & Business Media, 2012.
Cited on page 176
- [EH10] Herbert Edelsbrunner and John Harer. *Computational topology: an introduction*. American Mathematical Soc., 2010.
Cited (2) on pages 31 and 71
- [FC17] R’emi Flamary and Nicolas Courty. Pot python optimal transport library, 2017.
Cited (2) on pages 159 and 177
- [FG10] Alessio Figalli and Nicola Gigli. A new transportation distance between non-negative measures, with applications to gradients flows with dirichlet boundary conditions. *Journal de mathématiques pures et appliquées*, 94(2):107–130, 2010.
Cited (14) on pages 52, 53, 64, 65, 67, 76, 78, 81, 83, 85, 88, 89, 90, and 101
- [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):707–738, 2015.
Cited on page 155
- [Fig10] Alessio Figalli. The optimal partial transport problem. *Archive for rational mechanics and analysis*, 195(2):533–560, 2010.
Cited on page 50
- [Fro92] Patrizio Frosini. Discrete computation of size functions. 1992.
Cited on page 22
- [FSV⁺19] Jean Feydy, Thibault Séjourné, François-Xavier Vialard, Shun-ichi Amari, Alain Trounev, and Gabriel Peyré. Interpolating between optimal transport and mmd using sinkhorn divergences. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 2681–2690, 2019.
Cited (3) on pages 55, 56, and 58
- [Gab72] Peter Gabriel. Unzerlegbare darstellungen i. *Manuscripta mathematica*, 6(1):71–103, 1972.
Cited on page 213
- [GBR⁺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(Mar):723–773, 2012.
Cited on page 58

- [Gen19] Aude Genevay. *Entropy-regularized optimal transport for machine learning*. PhD thesis, 2019. *Cited (2) on pages 56 and 174*
- [GPC18] Aude Genevay, Gabriel Peyre, and Marco Cuturi. Learning generative models with sinkhorn divergences. In *International Conference on Artificial Intelligence and Statistics*, pages 1608–1617, 2018. *Cited (2) on pages 55 and 58*
- [GPRS19] Thibaut Le Gouic, Quentin Paris, Philippe Rigollet, and Austin J Stromme. Fast convergence of empirical barycenters in alexandrov spaces and the wasserstein space. *arXiv preprint arXiv:1908.00828*, 2019. *Cited (2) on pages 50 and 182*
- [GUD15] GUDHI Project. *GUDHI User and Reference Manual*. GUDHI Editorial Board, 2015. *Cited on page 19*
- [Hal86] Marshall Hall. *Combinatorial theory (2nd ed)*. John Wiley & Sons, 1986. *Cited on page 73*
- [HKN19] Christoph D. Hofer, Roland Kwitt, and Marc Niethammer. Learning representations of persistence barcodes. *Journal of Machine Learning Research*, 20(126):1–45, 2019. *Cited on page 139*
- [HKNU17] Christoph Hofer, Roland Kwitt, Marc Niethammer, and Andreas Uhl. Deep learning with topological signatures. In *Advances in Neural Information Processing Systems*, pages 1634–1644, 2017. *Cited (3) on pages 140, 143, and 152*
- [HSW07] Jan-Martin Hertzsch, Rob Sturman, and Stephen Wiggins. Dna microarrays: design principles for maximizing ergodic, chaotic mixing. *Small*, 3(2):202–218, 2007. *Cited on page 141*
- [Kal83] O Kallenberg. *Random measures*, akademie-verlag, 1983. *Cited on page 221*
- [Kal17] Olav Kallenberg. *Random measures, theory and applications*. Springer, 2017. *Cited on page 151*
- [Kan42] L Kantorovich. On the transfer of masses (in russian). In *Doklady Akademii Nauk*, volume 37, pages 227–229, 1942. *Cited on page 45*

- [Kec95] A.S. Kechris. *Classical Descriptive Set Theory*. Graduate Texts in Mathematics. Springer-Verlag, 1995.
Cited on page 153
- [KFH17] Genki Kusano, Kenji Fukumizu, and Yasuaki Hiraoka. Kernel method for persistence diagrams via kernel embedding and weight factor. *The Journal of Machine Learning Research*, 18(1):6947–6987, 2017.
Cited on page 44
- [KH⁺09] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
Cited (2) on pages 9 and 20
- [KHF16] Genki Kusano, Yasuaki Hiraoka, and Kenji Fukumizu. Persistence weighted gaussian kernel for topological data analysis. In *International Conference on Machine Learning*, pages 2004–2013, 2016.
Cited (3) on pages 136, 140, and 142
- [KHN⁺15] Roland Kwitt, Stefan Huber, Marc Niethammer, Weili Lin, and Ulrich Bauer. Statistical topological data analysis—a kernel perspective. In *Advances in neural information processing systems*, pages 3070–3078, 2015.
Cited on page 69
- [KKM⁺16] Kristian Kersting, Nils Kriege, Christopher Morris, Petra Mutzel, and Marion Neumann. Benchmark data sets for graph kernels, 2016.
Cited (2) on pages 9 and 20
- [KMN17] Michael Kerber, Dmitriy Morozov, and Arnur Nigmatov. Geometry helps to compare persistence diagrams. *Journal of Experimental Algorithmics (JEA)*, 22(1):1–4, 2017.
Cited (2) on pages 71 and 118
- [KMV⁺16] Stanislav Kondratyev, Léonard Monsaingeon, Dmitry Vorotnikov, et al. A new optimal transport distance on the space of finite radon measures. *Advances in Differential Equations*, 21(11/12):1117–1164, 2016.
Cited on page 50
- [KVT19] Max Kontak, Jules Vidal, and Julien Tierny. Statistical parameter selection for clustering persistence diagrams. In *2019 IEEE/ACM HPC for Urgent Decision Making (UrgentHPC)*, pages 7–12. IEEE, 2019.
Cited on page 118

- [LCO18] Théo Lacombe, Marco Cuturi, and Steve Oudot. Large scale computation of means and clusters for persistence diagrams using optimal transport. In *Advances in Neural Information Processing Systems*, 2018. *Cited (4) on pages 19, 29, 71, and 109*
- [Léo14] Christian Léonard. A survey of the schrödinger problem and some of its connections with optimal transport. *Discrete & Continuous Dynamical Systems-A*, 34(4):1533, 2014. *Cited on page 54*
- [LGL16] Thibaut Le Gouic and Jean-Michel Loubes. Existence and consistency of wasserstein barycenters. *Probability Theory and Related Fields*, pages 1–17, 2016. *Cited (4) on pages 50, 93, 94, and 96*
- [Luc10] Yves Lucet. What shape is your conjugate? a survey of computational convex analysis and its applications. *SIAM review*, 52(3):505–542, 2010. *Cited (2) on pages 120 and 125*
- [LY18] Tam Le and Makoto Yamada. Persistence Fisher kernel: a Riemannian manifold kernel for persistence diagrams. In *Advances in Neural Information Processing Systems*, pages 10027–10038, 2018. *Cited (4) on pages 44, 140, 141, and 142*
- [MKPY16] Nikolay Makarenko, Maksat Kalimoldayev, Ivan Pak, and Ainur Yessenaliyeva. Texture recognition by the methods of topological data analysis. *Open Engineering*, 6(1), 2016. *Cited on page 127*
- [MMH11] Yuriy Mileyko, Sayan Mukherjee, and John Harer. Probability measures on the space of persistence diagrams. *Inverse Problems*, 27(12):124007, 2011. *Cited (6) on pages 38, 39, 41, 42, 68, and 96*
- [Mon84] Gaspard Monge. *Mémoire sur la théorie des déblais et des remblais*. Histoire de l’Académie Royale des Sciences de Paris, 1784. *Cited on page 45*
- [Mor40] Marston Morse. Rank and span in functional topology. *Annals of Mathematics*, pages 419–454, 1940. *Cited on page 22*
- [Mor65] Jean-Jacques Moreau. Proximité et dualité dans un espace hilbertien. *Bulletin de la Société mathématique de France*, 93:273–299, 1965. *Cited on page 120*

- [Mun84] James R Munkres. *Elements of algebraic topology*, volume 2. Addison-Wesley Menlo Park, 1984.
Cited (4) on pages 12, 22, 203, and 206
- [Nes05] Yu Nesterov. Smooth minimization of non-smooth functions. *Mathematical programming*, 103(1):127–152, 2005.
Cited on page 162
- [Nie11] Lance Nielsen. Weak convergence and banach space-valued functions: improving the stability theory of feynman’s operational calculi. *Mathematical Physics, Analysis and Geometry*, 14(4):279–294, 2011.
Cited on page 135
- [OHK18] Ippei Obayashi, Yasuaki Hiraoka, and Masao Kimura. Persistence diagrams with linear machine learning models. *Journal of Applied and Computational Topology*, 1(3-4):421–449, 2018.
Cited on page 127
- [Oud15] Steve Y Oudot. *Persistence theory: from quiver representations to data analysis*, volume 209. American Mathematical Society, 2015.
Cited on page 31
- [OY19] Ippei Obayashi and Michio Yoshiwaki. Field choice problem in persistent homology. *arXiv preprint arXiv:1911.11350*, 2019.
Cited on page 206
- [PC17] Gabriel Peyré and Marco Cuturi. *Computational Optimal Transport*. Number 2017-86. December 2017.
Cited (7) on pages 42, 45, 56, 71, 117, 120, and 173
- [Per74] Michael D Perlman. Jensen’s inequality for a convex vector-valued function on an infinite-dimensional space. *Journal of Multivariate Analysis*, 4(1):52–65, 1974. *Cited on page 154*
- [PMK19] Jose A Perea, Elizabeth Munch, and Firas A Khasawneh. Approximating continuous functions on persistence diagrams using template functions. *arXiv preprint arXiv:1902.07190*, 2019.
Cited on page 76
- [Pra07] Aldo Pratelli. On the equality between monge’s infimum and kantorovich’s minimum in optimal mass transportation. In *Annales de l’Institut Henri Poincaré (B) Probability and Statistics*, volume 43, pages 1–13. Elsevier, 2007.
Cited on page 48

- [PWZ⁺17] Shirui Pan, Jia Wu, Xingquan Zhu, Guodong Long, and Chengqi Zhang. Task sensitive feature exploration and learning for multitask graph classification. *IEEE transactions on cybernetics*, 47(3):744–758, 2017. *Cited (2) on pages 9 and 20*
- [RHBK15] Jan Reininghaus, Stefan Huber, Ulrich Bauer, and Roland Kwitt. A stable multi-scale kernel for topological machine learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4741–4748, 2015. *Cited (5) on pages 44, 136, 140, 142, and 152*
- [Rob99] Vanessa Robins. Towards computing homology from finite approximations. In *Topology proceedings*, volume 24, pages 503–532, 1999. *Cited on page 22*
- [RSV18] Frédéric Le Roux, Sobhan Seyfaddini, and Claude Viterbo. Barcodes and area-preserving homeomorphisms. *arXiv preprint arXiv:1810.03139*, 2018. *Cited on page 170*
- [RTC17] Aaditya Ramdas, Nicolás García Trillos, and Marco Cuturi. On wasserstein two-sample testing and related families of nonparametric tests. *Entropy*, 19(2):47, 2017. *Cited (2) on pages 55 and 58*
- [San15] Filippo Santambrogio. Optimal transport for applied mathematicians. *Birkäuser, NY*, 2015. *Cited (4) on pages 45, 72, 78, and 120*
- [Sch31] Erwin Schrödinger. *Über die umkehrung der naturgesetze*. Verlag der Akademie der Wissenschaften in Kommission bei Walter De Gruyter u . . . , 1931. *Cited on page 54*
- [Sch03] Alexander Schrijver. *Combinatorial optimization: polyhedra and efficiency*, volume 24. Springer Science & Business Media, 2003. *Cited (2) on pages 104 and 105*
- [SDGP⁺15] Justin Solomon, Fernando De Goes, Gabriel Peyré, Marco Cuturi, Adrian Butscher, Andy Nguyen, Tao Du, and Leonidas Guibas. Convolutional wasserstein distances: Efficient optimal transportation on geometric domains. *ACM Transactions on Graphics (TOG)*, 34(4):66, 2015. *Cited (2) on pages 50 and 122*

- [SFV⁺19] Thibault Séjourné, Jean Feydy, François-Xavier Vialard, Alain Trounev, and Gabriel Peyré. Sinkhorn divergences for unbalanced optimal transport. *arXiv preprint arXiv:1910.12958*, 2019. *Cited (2) on pages 55 and 58*
- [SKC18] Don Sheehy, Oliver Kisielius, and Nicholas J Cavanna. Computing the shift-invariant bottleneck distance for persistence diagrams. In *CCCG*, pages 78–84, 2018. *Cited on page 171*
- [SP04] Robert W Sumner and Jovan Popović. Deformation transfer for triangle meshes. In *ACM Transactions on Graphics (TOG)*, volume 23, pages 399–405. ACM, 2004. *Cited on page 128*
- [SP18] Shashank Singh and Barnabás Póczos. Minimax distribution estimation in wasserstein distance. *arXiv preprint arXiv:1802.08855*, 2018. *Cited on page 157*
- [STNR⁺18] Anirudh Som, Kowshik Thopalli, Karthikeyan Natesan Ramamurthy, Vinay Venkataraman, Ankita Shukla, and Pavan Turaga. Perturbation robust representations of topological persistence diagrams. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 617–635, 2018. *Cited on page 136*
- [TMMH14] Katharine Turner, Yuriy Mileyko, Sayan Mukherjee, and John Harer. Fréchet means for distributions of persistence diagrams. *Discrete & Computational Geometry*, 52(1):44–70, 2014. *Cited (13) on pages 38, 40, 41, 42, 71, 110, 112, 113, 115, 116, 127, 128, and 129*
- [Tur13] Katharine Turner. Means and medians of sets of persistence diagrams. *arXiv preprint arXiv:1307.8300*, 2013. *Cited (2) on pages 42 and 96*
- [TVH18] Quoc Hoan Tran, Van Tuan Vo, and Yoshihiko Hasegawa. Scale-variant topological information for characterizing complex networks. *arXiv preprint arXiv:1811.03573*, 2018. *Cited (2) on pages 143 and 145*
- [Ume17] Yuhei Umeda. Time series classification via topological data analysis. *Information and Media Technologies*, 12:228–239, 2017. *Cited on page 134*

- [VBT19] Jules Vidal, Joseph Budin, and Julien Tierny. Progressive wasserstein barycenters of persistence diagrams. *IEEE transactions on visualization and computer graphics*, 26(1):151–161, 2019. *Cited on page 118*
- [Vil08] Cédric Villani. *Optimal transport: old and new*, volume 338. Springer Science & Business Media, 2008. *Cited (9) on pages 45, 47, 48, 49, 67, 76, 87, 89, and 90*
- [VZ17] Saurabh Verma and Zhi-Li Zhang. Hunt for the unique, stable, sparse and fast feature learning on graphs. In *Advances in Neural Information Processing Systems*, pages 88–98, 2017. *Cited (2) on pages 143 and 145*
- [Wag19] Alexander Wagner. Nonembeddability of persistence diagrams with $p > 2$ wasserstein metric. *arXiv preprint arXiv:1910.13935*, 2019. *Cited (2) on pages 16 and 27*
- [WCP13] Weiran Wang and Miguel A Carreira-Perpinán. Projection onto the probability simplex: An efficient algorithm with a simple proof, and an application. *arXiv preprint arXiv:1309.1541*, 2013. *Cited on page 165*
- [XC19] Zhang Xinyi and Lihui Chen. Capsule graph neural network. In *International Conference on Learning Representations*, 2019. *Cited (2) on pages 143 and 145*
- [XHLJ19] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? In *ICLR 2019*, 2019. *Cited (2) on pages 143 and 145*
- [YV15] Pinar Yanardag and S.V.N. Vishwanathan. Deep graph kernels. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '15*, pages 1365–1374, New York, NY, USA, 2015. ACM. *Cited (3) on pages 9, 20, and 143*
- [ZC05] Afra Zomorodian and Gunnar Carlsson. Computing persistent homology. *Discrete & Computational Geometry*, 33(2):249–274, 2005. *Cited on page 22*
- [ZKR⁺17] Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Ruslan Salakhutdinov, and Alexander Smola. Deep sets. In *Advances in Neural Information Processing Systems*, pages 3391–3401, 2017. *Cited on page 139*

- [ZWX⁺18] Zhen Zhang, Mianzhi Wang, Yijian Xiang, Yan Huang, and Arye Nehorai. RetGK: Graph Kernels based on Return Probabilities of Random Walks. In *Advances in Neural Information Processing Systems*, pages 3968–3978, 2018. *Cited (2) on pages 143 and 145*
- [ZZJS16] Matthias Zeppelzauer, Bartosz Zieliński, Mateusz Juda, and Markus Seidl. Topological descriptors for 3d surface analysis. In *International Workshop on Computational Topology in Image Context*, pages 77–87. Springer, 2016. *Cited on page 127*

INDEX

- φ -divergence, 51
- c -concave, 77
- c -transform, 77
- p -th distance, 36
- Čech, 167
- Čech complex, 208
- Čech diagrams, 216

- absolutely continuous, 220
- admissible plans, 64
- admissible transport plans, 52
- angle, 222

- B-Munkres algorithm, 112
- barycenter, 41
- birth times, 23
- bottleneck distance, 24, 33
- boundaries, 204
- boundary operator, 203

- chain complex, 203
- convex indicator function, 52
- convexity, 95
- cost, 33
- critical values, 23, 164
- Csiszár-divergence, 51
- curvature, 41
- cycles, 203
- cyclically monotone, 77

- death times, 23
- Deep Sets, 137

- direct sum of modules, 211
- divergence, 58

- energy functional, 41, 49, 92
- entropic regularization, 54, 112, 116, 145
- entropy function, 51
- envelope theorem, 159, 166
- essential part, 34
- Eulerian, 116
- expectation, 20
- expected persistence diagram, 43, 146
- expected persistence diagrams, 145

- filtration, 23, 32, 164, 206
- finite-dimensional setting, 56
- Fréchet mean, 41, 49, 92

- geodesic, 41, 221
- geodesic space, 221
- geometric realization, 205
- graph classification, 140
- Gromov-Hausdorff distance, 217
- ground metric, 59

- Hausdorff distance, 216
- Heat Kernel Signatures, 140
- histograms, 57
- homology, 201
- homology group, 204
- homotopic, 204

- homotopy type, 204
- integer programming, 96
- interleaving, 214
- interleaving distance, 216
- interval module, 32, 210
- Kantorovich, 47
- Kullback-Leibler divergence, 52, 55
- Lagrangian, 116
- length, 221
- length space, 221
- linear representation, 44, 132
- marginal constraints, 48
- maximum mean discrepancy, 57
- Monge problem, 46
- multiplicity, 212
- negentropy, 54, 117
- nerve theorem, 208
- non-negatively curved Alexandrov space, 222
- offsets filtration, 207
- optimal matching, 48
- optimal transport cost, 47, 48
- optimal transport map, 47
- optimal transport plan, 48
- orientation of a simplex, 202
- partial matching, 33
- partial transport, 25
- path, 221
- permutation invariant, 137, 157
- permutation invariant operator, 136
- persistence diagram, 23, 32, 201, 212
- persistence measures, 64
- persistence module, 32, 208
- persistence theory, 206
- point measure, 32, 38
- point measures, 220
- point transformation, 136
- Polish metric space, 219
- product measure, 54
- pseudo-metrics, 37
- push-forward, 45, 168
- q-tame, 35, 212
- quadratic cost, 48
- quantization, 29, 145, 155
- Radon measures, 219
- Radon-Nikodym, 51
- random persistence diagram, 43
- representation function, 136
- Rips complex, 208
- sample, 20
- separable cost, 36, 112
- shift-invariant, 145
- simplex, 201
- simplicial complex, 201
- simplicial filtration, 213
- simplicial homology, 201, 202
- singular, 220
- singular simplex, 205
- Sinkhorn cost, 117
- Sinkhorn divergence, 57
- Sinkhorn map, 56, 117
- space of persistence diagrams, 38, 59
- space of persistence measures, 59
- Stability theorem, 35
- sublevel sets, 206
- superdifferential, 77
- supervised setting, 135
- symplectic topology, 168
- tight, 219
- topological data analysis, 21
- torsion, 204
- total persistence, 38, 59, 163, 171

total variation, 52
transport map, 46
transport plans, 47
transportation polytope, 48, 54

unbalanced optimal transport, 51

vague topology, 219
variance, 20
vectorization, 26, 44, 157
vertices, 201

Wasserstein barycenter, 49
Wasserstein distance, 48
Wasserstein space, 49
Weak topology, 219
weight function, 136

APPENDIX A

COMPLEMENTARY NOTIONS OF HOMOLOGY AND PERSISTENCE THEORY

This chapter recalls the algebraic root of topological data analysis: *homology* theory, which gives a formal meaning to the notions of “topological properties”. It then presents (one-dimensional) persistence theory, the modern bridge between homology and data science, and its main tool: the *persistence diagram*.

A.1 Homology theory

For the sakes of simplicity and concision, we focus in this section on *simplicial homology*, which is the version of homology used in applications. We briefly mention at the end of this section an extension to more general objects. We refer the interested reader to [Mun84] for more details.

Simplicial complexes

Simplicial homology is the formal way to describe the topology of a simplicial complex, a generalization of the notion of graph that can be conveniently encoded in a computer and used to make computations in applications.

DEFINITION A.1. *Let $E = \{v_0 \dots v_N\}$ be a finite set. A simplex on E is an element of $\mathcal{P}(E)$, whose elements are called its vertices. A simplicial complex X on E is a (non-empty) subset of $\mathcal{P}(E)$ that satisfies the*




Figure A.1: An example of simplicial complex.

following:

$$\forall \sigma \in X, \forall \tau \subset \sigma \text{ with } \tau \neq \emptyset, \text{ we have } \tau \in X.$$

The dimension of a simplex σ is $\dim(\sigma) := \text{card}(\sigma) - 1$. Simplices of dimension 0 will be called points or vertices, those of dimension 1 are called edges, dimension 2 and above are generally called faces. The dimension of a simplicial complex is the largest dimension of one of the simplices it contains. The k -simplices of X are the simplices of dimension k in X , and we write X_k for the set of k -simplices of X .

Note that simplicial complexes are purely combinatorial structures: they do not need to be embedded in some (Euclidean) space.

We now introduce the notion of *orientation of a simplex*. Intuitively, it formalizes the idea that an edge $\{a, b\}$ can be oriented either “from a to b ”, denoted by $[a, b]$, or “from b to a ”, denoted by $[b, a]$.

DEFINITION A.2 (Orientation). *Let $\sigma = \{v_0 \dots v_k\}$ be a k -simplex. Two orderings $(v_i)_{i=1}^k$ of σ are equivalent if they differ from one another by an even permutation. It defines two equivalence classes, called the orientations of σ . We will write $\sigma = [v_0 \dots v_k]$ to specify an orientation of σ .*

From now on, we will implicitly assume that our simplices are oriented.

EXAMPLE A.3. *The complex represented in Figure A.1 is defined by*

$$\begin{aligned} X = \{ \{a\}, \{b\}, \{c\}, \{d\}, & \quad \text{simplices of dim 0} \\ \{ab\}, \{ad\}, \{bc\}, \{bd\}, \{cd\} & \quad \text{simplices of dim 1} \\ \{abd\} \} & \quad \text{simplices of dim 2.} \end{aligned}$$

Simplicial homology

Computing the *simplicial homology* of a simplicial complex consists, roughly speaking, in encoding its topological features. Informally, topological features are connected components (features of “dimension 0”), loops (dimension 1), cavities (dimension 2), and so on... To formalize this intuition, we need to put an algebraic structure on our simplicial complex that will allow us, both in theory and in practice, to identify these topological features.

Let X be a (finite) simplicial complex. Let \mathbb{Z} denote the ring of integers.

DEFINITION A.4 (*k*-chains). *The space of k-chains of a simplicial complex X is the free abelian group generated by the set of k-simplices $X_k = \{\sigma_1 \dots \sigma_{|X_k|}\}$:*

$$C_k(X, \mathbb{Z}) := \left\{ \sum_{i=1}^{|X_k|} \alpha_i \sigma_i : \alpha_i \in \mathbb{Z} \right\}.$$

boundary operator. Now that we have a group structure on the $(X_k)_k$, we can build morphisms between these.

DEFINITION A.5. *Given a k-simplex $[v_0 \dots v_k] \in X_k$, we denote by $[v_0 \dots, \widehat{v}_j, \dots, v_k] \in X_{k-1}$ the (k-1)-simplex (which is a face of $[v_0 \dots v_k]$) where the vertex v_j has been removed.*

The boundary operator ∂_k for $k \geq 1$ over *k*-dimensional chains is defined as the morphism:

$$\begin{aligned} \partial_k : C_k(X, \mathbb{Z}) &\rightarrow C_{k-1}(X, \mathbb{Z}) \\ \sigma = [v_0 \dots v_k] &\mapsto \sum_{j=0}^k (-1)^j \underbrace{[v_0 \dots \widehat{v}_j, \dots, v_k]}_{\in X_{k-1}} \\ (\sigma + \lambda \sigma') &\mapsto \partial_k \sigma + \lambda \partial_k \sigma' \end{aligned}$$

By convention, $\partial_0 = 0$.

Note that $\forall k > 0$ we have $\partial_{k-1} \circ \partial_k = 0$, that is, the boundary of a boundary is zero.

Homology groups. We have the following sequence¹ of applications:

$$C_k(X, \mathbb{Z}) \xrightarrow{\partial_k} C_{k-1}(X, \mathbb{Z}) \xrightarrow{\partial_{k-1}} C_{k-2}(X, \mathbb{Z}) \xrightarrow{\partial_{k-2}} \dots$$

¹Called a *chain complex*, as if $A \xrightarrow{f} B \xrightarrow{g} C$, we have $\text{Im}(f) \subset \text{ker}(g)$.

We define:

- *k-cycles* : $Z_k := \ker(\partial_k) = \{\sigma \in C_k, \partial_k \sigma = 0\}$
- *k-boundaries* : $B_k := \text{Im}(\partial_{k+1}) = \{\partial_{k+1} \sigma, \sigma \in C_{k+1}\}$

As $\partial_{k-1} \circ \partial_k = 0$, we know that $\text{Im}(\partial_{k+1}) \subset \ker(\partial_k)$, which states that boundaries of $(k+1)$ -simplices are k -cycles. It allows us to define the following quotient.

DEFINITION A.6. *The k -th homology group of X is*

$$H_k(X, \mathbb{Z}) := \frac{\ker(\partial_k)}{\text{Im}(\partial_{k+1})}. \quad (\text{A.1})$$

This definition can be interpreted in the following way: the homology group of a complex X represents the “space of its cycles modulo its boundaries”.

The collection of the homology groups of X is denoted by $H_*(X, \mathbb{Z})$.

REMARK A.7. *In practice, simplicial homology is computed by replacing the ring \mathbb{Z} by a finite field \mathbb{K} such as $\mathbb{Z}/p\mathbb{Z}$ for some prime number p (typically $p = 2$). It follows that $C_k(X, \mathbb{K})$ and thus $H_k(X, \mathbb{K})$ have a finite-dimensional vector space structure. In this context, the dimension β_k of $H_k(X, \mathbb{K})$ is called the k -th Betti number of X (for the field of coefficients \mathbb{K}). A basis $v_1 \dots v_{\beta_k}$ of $H_k(X, \mathbb{K})$, which are (representing) linear combinations of elements in X , can be understood as a maximal family of independent topological features of dimension k . In particular, β_k encodes the number of such topological features.*

From a theoretical perspective, the choice of the field of coefficients \mathbb{K} has some consequence on the computation of $H_k(X, \mathbb{K})$ only when $H_k(X, \mathbb{Z})$ has torsion, see [Mun84] for more details. In applications however, it seems that the choice of \mathbb{K} has barely any impact [OY19].

Invariance properties. We briefly discuss here why homology is a relevant descriptor to describe the topology of structured objects. First, let us introduce a notion of equivalence between topological spaces.

DEFINITION A.8. *Let \mathcal{X} and \mathcal{Y} be two topological spaces. Two continuous maps $u, v : \mathcal{X} \rightarrow \mathcal{Y}$ are homotopic if there exists $\varphi : [0, 1] \times \mathcal{X} \rightarrow \mathcal{Y}$ continuous, such that $\varphi(0, \cdot) = u$ and $\varphi(1, \cdot) = v$. This will be denoted by $u \sim v$.*

Let \mathcal{X} and \mathcal{Y} be two topological spaces. They are said to have the same homotopy type if there exist $f : \mathcal{X} \rightarrow \mathcal{Y}$ and $g : \mathcal{Y} \rightarrow \mathcal{X}$ such as $f \circ g \sim \text{id}_{\mathcal{Y}}$, $g \circ f \sim \text{id}_{\mathcal{X}}$. We will denote this by $\mathcal{X} \sim \mathcal{Y}$.

Let X be a (finite) simplicial complex. A *geometric realization* of X is a map $\psi : X \rightarrow \mathbb{R}^{2n+1}$ such that for any simplex $\sigma \in X$, with $\sigma = \{v_1 \dots v_k\}$, one has $\psi(\sigma) = \text{conv}(\{\psi(v_1) \dots \psi(v_k)\})$ where the $(\psi(v_i))_i$ must be affinely independent and $\text{conv}(A)$ denotes the convex hull of a set A , and such that for any $\sigma_1, \sigma_2 \in X$, one has $\psi(\sigma_1) \cap \psi(\sigma_2) = \psi(\sigma_1 \cap \sigma_2)$ (with $\psi(\emptyset) = \emptyset$). A simplicial complex X of dimension n always admits a geometric realization in the Euclidean space \mathbb{R}^{2n+1} (by Whitney's Embedding theorem).

The important result to note is the following one:

THEOREM. *Let X, Y be two simplicial complexes, and $r(X), r(Y)$ be two geometric realizations of those. If $r(X) \sim r(Y)$ then $H_*(X, \mathbb{K}) = H_*(Y, \mathbb{K})$.*

This theorem can be rephrased by stating that two complexes that have isomorphic homotopy type have the same homology groups, or equivalently that homology is a homotopy invariant.

Singular homology. We end this section by mentioning that homology theory can be extended to a wider class of objects than just simplicial complexes. Namely, given a topological space \mathcal{X} , we can say that $\Sigma \subset \mathcal{X}$ is a *singular simplex* of X of dimension k if it is the image of the geometric realization of a (oriented) k -simplex $\text{conv}(\{v_0 \dots v_k\})$ by a continuous map (Figure A.2). All the definitions and properties stated above can be adapted to general topological spaces: the space of k -chains is the free abelian group generated by the (oriented) singular simplices of X , and all the subsequent definitions (boundary operator, cycles, boundaries, homology groups, homotopy invariance) extend straightforwardly. Note that the (simplicial) homology groups of X and the (singular) homology groups of a geometric realization $r(X)$ do coincide, making singular homology a satisfactory extension of simplicial homology to general topological spaces. Note that singular homology is also a homotopy invariant.

A.2 Filtrations and persistence modules

Although simplicial homology allows us to describe the topological features of a given object, it cannot be used faithfully in statistical and machine learning applications, for several reasons:




Figure A.2: (Left) A geometric realization Δ of the 2-simplex in \mathbb{R}^2 . (Right) A torus X and a singular simplex (the image of Δ by a continuous map σ).

- The topology of a simplicial complex is based on a purely combinatorial representation of the object, without (explicitly) taking the extra information available (such as the geometry of the ambient space) into account.
- The homology groups of a point cloud $P \subset \mathbb{R}^d$ with n points will simply be given by $H_0(P, \mathbb{Z}) = \mathbb{Z}^n$ and $H_k(P, \mathbb{Z}) = 0$ for $k \geq 1$ (n connected components and nothing else); even if the points are perfectly distributed on, say, a circle.
- The information encoded by homology might not be quantitative enough. Although $H_1(X, \mathbb{Z})$ might account for the presence of loops, for instance, we do not have access to some quantitative information about these loops: are there some loops more significant, or more reliable, than others?

To improve on these, we will rely on *persistence theory*: namely we will study the evolution of the homology groups of an object at various scales using a *filtration*.

Preliminary definitions

DEFINITION A.9. Let \mathcal{X} be a topological space and $T \subset \mathbb{R}$. A filtration $(\mathcal{F}_t)_{t \in T}$ is a family of topological spaces such that

$$\forall (s, t) \in T^2, \quad s \leq t \Rightarrow \mathcal{F}_s \subset \mathcal{F}_t.$$




Figure A.3: Example of Čech and Rips complexes built on a point cloud.

A standard way to build filtration is to look at the sublevel sets of a map $f : \mathcal{X} \rightarrow \mathbb{R}$, i.e. the subsets of \mathcal{X} of the form

$$\mathcal{F}_t := f^{-1}((-\infty, t]).$$

EXAMPLE A.10. Let $P = \{x_1 \dots x_n\}$ be a point cloud in \mathbb{R}^d . The offsets filtration on P is the filtration described by the map

$$f = \text{dist}_P : x \mapsto \min_{1 \leq i \leq n} \|x - x_i\|_2.$$

Here, one has

$$\mathcal{F}_t = \bigcup_{i=1}^n B(x_i, t),$$

where $B(x_i, t)$ is the Euclidean ball centered in x_i of radius t for $t \geq 0$, set to be the empty set if $t < 0$. See Figure A.3 for an example. Note that in practice, computing the homology groups of the space \mathcal{F}_t directly is not handy as algorithms take simplicial complexes as input. In this particular

case however, one has two natural options to build a simplicial complex based on \mathcal{F}_t :

- The Čech complex $C(P, t)$ of parameter t is given, for $t \geq 0$, by

$$\{x_{i_1} \dots x_{i_k}\} = \sigma \in C(P, t) \Leftrightarrow \bigcap_{j=0}^k B(x_{i_j}, t) \neq \emptyset,$$

that is the balls have a common intersection. The nerve theorem ensures that $C(P, t)$ and \mathcal{F}_t have the same homotopy type and thus the same homology groups (see Appendix A). Therefore, working in practice with the simplicial complexes $(C(P, t))_t$ allows us to exactly recover the topology of $(\mathcal{F}_t)_t$.

- Similarly, the Rips complex $R(P, t)$ of parameter t is given, for $t \geq 0$, by

$$\{x_{i_1} \dots x_{i_k}\} = \sigma \in R(P, t) \Leftrightarrow \forall j, j', B(x_{i_j}, t) \cap B(x_{i_{j'}}, t) \neq \emptyset,$$

that is the balls have pairwise intersections. Note that we have for all $t \geq 0$:

$$C(P, t) \subset R(P, t) \subset C(P, 2t). \quad (\text{A.2})$$

Furthermore, $R(P, t)$ is much simpler to compute than $C(P, t)$ as one has to check if the balls have pairwise intersections instead of common intersections.

Persistence modules

Now, fix a field \mathbb{K} . It turns out that $(H_*(\mathcal{F}_t, \mathbb{K}))_t$ is not a family of independent vector spaces but—thanks to the increasing property $s \leq t \Rightarrow \mathcal{F}_s \subset \mathcal{F}_t$ —has much more structure, which is encoded as a *persistence module*.

DEFINITION A.11. A persistence module over $T \subset \mathbb{R}$ is a family $\mathbb{V} = (V_t)_{t \in T}$ of \mathbb{K} -vector spaces, endowed with linear applications $v_t^{t'} : V_t \rightarrow V_{t'}$ such that:

$$\begin{aligned} \forall t \in T, \quad v_t^t &= id \\ \forall t \leq t' \leq t'' \in T, \quad v_{t'}^{t''} \circ v_t^{t'} &= v_t^{t''}. \end{aligned}$$

Note that the linear applications are generally made implicit in the notation, that is we will write $\mathbb{V} = (V_t)_t$ to denote a persistence module over a family of \mathbb{K} -vector spaces $(V_t)_t$ endowed with maps $(v_t^{t'})_{t \leq t'}$.




Figure A.4: Global picture of the algebraic pipeline behind the construction of persistence modules.

Persistence module of a simplicial filtration. Let X, Y, Z be three simplicial complexes such that $X \subset Y \subset Z$. Let i_X^Y and i_Y^Z denote the inclusion maps between (X, Y) and (Y, Z) respectively, along with $i_X^Z = i_Y^Z \circ i_X^Y$, the inclusion between X and Z . Let also H_X, H_Y, H_Z denote their respective homology groups. The inclusion maps are naturally turned into linear maps $v_{H_X}^{H_Y}, v_{H_Y}^{H_Z}, v_{H_X}^{H_Z}$ respectively, and one has $v_{H_Y}^{H_Z} \circ v_{H_X}^{H_Y} = v_{H_X}^{H_Z}$, see Figure A.5.

Thus, if we have $\mathcal{F} = (F_t)_t$ a filtration of simplicial complexes (called a simplicial filtration), we have:

- $\forall t \in T$, we define $V_t := H_*(F_t, \mathbb{K})$
- $\forall t \leq t'$, let $v_t^{t'}$ be the linear mapping induced by the canonical inclusion $F_t \hookrightarrow F_{t'}$.

This defines the (canonical) persistence module over \mathcal{F} . Without introducing the terminology of category theory, which is out of the scope covered by this manuscript, it can be useful to have the following global picture (Figure A.4) in mind: a real number s is turned into a topological space \mathcal{F}_s which is then turned into a vector space $H_*(\mathcal{F}_s, \mathbb{K})$, while inequalities $s \leq t$ are turned into inclusions $\mathcal{F}_s \subset \mathcal{F}_t$ which are then turned into linear maps $v_s^t : H_*(\mathcal{F}_s, \mathbb{K}) \rightarrow H_*(\mathcal{F}_t, \mathbb{K})$.

EXAMPLE A.12. Figure A.5 gives an example of an increasing sequence of simplicial complexes on which we derive the corresponding persistence module, that is $H_0(F_t) \oplus H_1(F_t), t \in \{1 \dots 6\}$. Linear applications $v_t^{t+1} :$

$H_0(F_t) \oplus H_1(F_t) \rightarrow H_0(F_{t+1}) \oplus H_1(F_{t+1})$ are represented as matrices at each step.




Figure A.5: An example of filtration over a simplicial complex and the corresponding persistence module.

A.3 Persistence diagrams

Decompositions of persistence modules

Fix $T \subset \mathbb{R}$ and a field \mathbb{K} .

DEFINITION A.13. An interval of T is a subset $I \subset T$ such as:

$$\forall t \leq t' \leq t'' \in T, \quad t, t'' \in I \Rightarrow t' \in I$$

An interval module over an interval $I \subset T$ is a persistence module \mathbb{V} defined by:

- $V_t = \mathbb{K}$ if $t \in I$, $V_t = \{0\}$ otherwise.
- $\forall t \leq t' \in T, v_t^{t'} = \text{id}_{\mathbb{K}}$ if $t, t' \in I$, $v_t^{t'} = 0$ otherwise.

For $I = [b, d]$, the corresponding interval module is denoted by $\mathbb{I}_{[b,d]}$. It can be described as

$$\mathbb{I}_{[b,d]} := \underbrace{\{0\} \xrightarrow{0} \dots \xrightarrow{0} \{0\}}_{t < b} \xrightarrow{0} \underbrace{\mathbb{K} \xrightarrow{\text{id}} \dots \xrightarrow{\text{id}} \mathbb{K}}_{b \leq t \leq d} \xrightarrow{0} \underbrace{\{0\} \xrightarrow{0} \dots \xrightarrow{0} \{0\}}_{d < t}$$

Analogous definitions and notations stand for $I = (a, b], (a, b), [a, b)$.

Interpretation. Interval modules are the “base blocks” of persistence modules, intuitively in the same way as prime numbers are the base block of natural integers, or (maybe more precisely) as real lines are the base blocks of real vector spaces. Just as any real vector space can be decomposed as a unique (up to isomorphism) direct sum of real lines, we will see that (under mild assumption), any persistence module can be uniquely (up to isomorphism) decomposed as a direct sum (see below) of interval modules.

DEFINITION A.14. Given $\mathbb{V} = (V_t)_t, \mathbb{W} = (W_t)_t$ two persistence modules with linear applications $(v_t^{t'})_{t,t'}, (w_t^{t'})_{t,t'}$, we define the direct sum of modules $\mathbb{V} \oplus \mathbb{W}$ by:

- $\mathbb{V} \oplus \mathbb{W} = (V_t \oplus W_t)_t$
- The corresponding linear applications are denoted by:

$$(v \oplus w)_t^{t'} : V_t \oplus W_t \rightarrow V_{t'} \oplus W_{t'}$$

$$(x, y) \mapsto (v_t^{t'}(x), w_t^{t'}(y))$$

This definition extends naturally to a family of persistence modules $(\mathbb{V}_j)_{j \in J}$, denoted by $\mathbb{V} := \bigoplus_{j \in J} \mathbb{V}_j$.

Note that $\mathbb{V} \oplus \mathbb{W}$ is also a persistence module (so is the sum over a family).

THEOREM A.15 ([CDSGO16, Theorem 2.8]). A persistence module \mathbb{V} can be decomposed as a direct sum of interval modules, written as:

$$\mathbb{V} \simeq \bigoplus \mathbb{I}_{[b_j, d_j]},$$

in the following sufficient cases:

1. If T is finite [Gab72].
2. When all the vector spaces V_t are finite-dimensional [CB15].

When it exists, the decomposition is unique (up to isomorphism and ordering of the terms).

When we have such a decomposition, the knowledge of the intervals $([b_j, d_j])_j$ gives a complete description of the structure of \mathbb{V} . This information can be encoded as a (multi-)set² of points, called the *persistence*

²That is, a set where elements can be repeated. The number of times an element x appears in a multiset is called its *multiplicity*.

diagram $\text{Dgm}(\mathbb{V})$ of the persistence module, supported on

$$\overline{\mathbb{R}}_{\geq}^2 := \{(b, d) \in \overline{\mathbb{R}}^2, \quad b \leq d\}, \quad (\text{A.3})$$

where $\overline{\mathbb{R}}$ is the extended real line $\mathbb{R} \cup \{\pm\infty\}$. Each point in the persistence diagram accounts for an interval module in the decomposition of \mathbb{V} .

It might happen that \mathbb{V} is not decomposable. However, one can still define a persistence diagram of \mathbb{V} assuming that \mathbb{V} is *q-tame*.

DEFINITION A.16. *A persistence module \mathbb{V} is said to be q-tame if $\forall t < t' \in T$, $\text{rank}(v_t^{t'})$ is finite.*

Extending the notion of persistence diagram for a persistence module \mathbb{V} that cannot be decomposed can be done using a measure-theoretic formalism. Intuitively, the idea is the following: when \mathbb{V} can be decomposed, the number of points of $\text{Dgm}(\mathbb{V})$ that fall into a rectangle $[a, b] \times [c, d] \subset \overline{\mathbb{R}}_{\geq}^2$ with $a < b \leq c < d$ corresponds to the number of intervals that start between a and b and end between c and d . It turns out to be equal to $\text{rank}(v_a^d) - \text{rank}(v_a^c) + \text{rank}(v_b^c) - \text{rank}(v_b^d)$.

Therefore, for a q-tame persistence module \mathbb{V} , let us define

$$\mu_{\mathbb{V}} : [a, b] \times [c, d] \mapsto \text{rank}(v_a^d) - \text{rank}(v_a^c) + \text{rank}(v_b^c) - \text{rank}(v_b^d),$$

which defines, modulo some technical details (see [CDSGO16, Chapter 3]), a measure on $\overline{\mathbb{R}}_{\geq}^2$. It turns out that this measure has some strong structure properties.

THEOREM A.17 ([CCSG+09a, CDSGO16]). *Let \mathbb{V} be a q-tame persistence module. Then $\mu_{\mathbb{V}}$ is a point measure, i.e. is integer-valued and supported on a locally finite set $X \subset \overline{\mathbb{R}}_{\geq}^2$. It can thus be written as*

$$\mu_{\mathbb{V}} = \sum_{x \in X} n_x \delta_x,$$

where δ_x denotes the Dirac mass located at $x \in \overline{\mathbb{R}}_{\geq}^2$, and $n_x \in \mathbb{N}$ denotes the multiplicity of x .

Furthermore, if \mathbb{V} can be decomposed as $\mathbb{V} = \bigoplus_j \mathbb{I}_{[b_j, d_j]}$, then $\mu_{\mathbb{V}} = \text{Dgm}(\mathbb{V})$, in sense that $\mu_{\mathbb{V}} = \sum_j \delta_{(b_j, d_j)}$.

This theorem is the first example in which adopting a measure-theoretic perspective provides a powerful framework to study topological descriptors. We will adopt this approach in chapters 3 and 4 and showcase its strength to deal with various statistical problems occurring in topological data analysis.

Computation

Let us now mention how persistence diagrams can be computed. In practice, algorithms take as input a *simplicial filtration*, that is a filtration over a simplicial complex K which satisfies:

- $T = \{0, 1, \dots, m\}$ (finite, so we have a decomposable persistence module).
- $K_0 = \emptyset, K_m = K$.
- $\forall t \in T, K_t$ is a finite simplicial complex, which is a sub-complex of K_{t+1} .

We can also assume that we actually only add one simplex at each step, that is $K_{t+1} \setminus K_t = \{\sigma_t\}$.

Persistence diagrams can be computed using Gaussian elimination, taking in addition the ordering of the simplices into account. The pipeline reads as follows

1. Let M be the matrix of the boundary operator ∂ .
2. Define

$$\text{low}(j) := \begin{cases} \max\{i \mid M_{ij} \neq 0\} \\ 0 \text{ if } M_{ij} = 0 \text{ for all } i \end{cases}$$

3. Use Gaussian elimination from the left to the right (see Algorithm 7).

Algorithm 7 Compute the barcode corresponding to a simplicial filtration

```

for  $j = 1 \dots m$  do
  while  $\exists i < j$  s.t.  $\text{low}(i) = \text{low}(j) \neq 0$  do
     $c_j \leftarrow c_j - \frac{M_{[\text{low}(i),j]}}{M_{[\text{low}(i),i]}} c_i$ 
  end while
end for
return The reduced matrix.

```

4. Building the diagram: after reduction each column with only 0 entries accounts for the birth of a cycle. Any other column j (with non 0) induces a boundary which trivializes (fills in) the cycle $i = \text{low}(j)$. Thus:

- Each column i with only 0 entries encodes the beginning of one interval module in the decomposition of the persistence module $\mathbb{I}_{[t(i),?]}$, where $t(i)$ denotes the insertion time of the corresponding simplex.
- Finite intervals are $\mathbb{I}_{[t(i),t(j)]}$ with $i = \text{low}(j)$
- Infinite intervals are $\mathbb{I}_{[t(i),+\infty)}$ when there is no j such that $i = \text{low}(j)$.

Note that in practice, we generally build one diagram for each homology dimension k .

EXAMPLE A.18. Let $\mathbb{K} = \mathbb{Z}/2\mathbb{Z}$. The pipeline is illustrated in Figure A.6. After reduction, the column of e_5 is empty, it means that inserting e_5 created a cycle (of dimension $\dim(e_5) = 1$, that is a loop). As $\text{low}(\sigma_1) = e_5$, it means that inserting σ_1 at $t = 6$ killed (filled) the cycle created by e_5 . It is recorded in the persistence diagram by adding a point with the corresponding insertion times $(5, 6)$. Similarly, e_4 created a cycle. As there is no σ such that $\text{low}(\sigma) = e_4$, the cycle did not get killed at the end of the filtration process, resulting in a point of coordinates $(4, +\infty)$ in the persistence diagram.

Thus, the decomposition of the persistence module is given by

$$\begin{aligned} H_0 &\simeq \mathbb{I}_{[1,+\infty)} \oplus \mathbb{I}_{[1,2)} \oplus \mathbb{I}_{[1,2)} \oplus \mathbb{I}_{[3,4)} \\ H_1 &\simeq \mathbb{I}_{[4,+\infty)} \oplus \mathbb{I}_{[5,6)}, \end{aligned}$$


and the corresponding persistence diagrams in homology dimension $k = 0$ and $k = 1$ reads, adopting the measure notation,

$$\begin{aligned} \text{Dgm}_0 &= \delta_{(1,+\infty)} + 2\delta_{(1,2)} + \delta_{(3,4)} \\ \text{Dgm}_1 &= \delta_{(4,+\infty)} + \delta_{(5,6)}. \end{aligned}$$

Interleaving and bottleneck distances

Using the bottleneck distance to compare persistence diagrams is motivated by algebraic considerations. The bridge between the bottleneck distance and persistence modules is made by re-interpreting this distance in terms of *interleaving*.

DEFINITION A.19. Let $\varepsilon > 0$. Let $\mathbb{V} = (V_t)_{t \in \mathbb{R}}$ and $\mathbb{W} = (W_t)_{t \in \mathbb{R}}$ be two persistence modules on \mathbb{R} . \mathbb{V} and \mathbb{W} are said to be ε -interleaved if there are two families of applications $\varphi = (\varphi_t)_t, \psi = (\psi_t)_t$ where $\varphi_t : V_t \rightarrow W_{t+\varepsilon}$



Leading to (matrix reduction):




Figure A.6: Pipeline to build a persistence diagram from a simplicial filtration. (Top) The simplicial filtration. (Middle) On the left, the matrix of the boundary operator ∂ . Note that simplices must be ordered with respect to their insertion time t . On the right, the matrix after reduction. (Bottom) The persistence diagram deduced from the reduced matrix.

and $\psi_t : W_t \rightarrow V_{t+\varepsilon}$ such that for all $t \leq t'$, the diagrams in Figure A.7 commute:




Figure A.7: Schematic representation of interleaving between persistence module.

Then, we define the interleaving distance as

$$d_i(\mathbb{V}, \mathbb{W}) := \inf \{ \varepsilon \geq 0 \mid \mathbb{V}, \mathbb{W} \text{ are } \varepsilon\text{-interleaved} \} \in \mathbb{R}_+ \cup \{+\infty\}.$$

The bottleneck distance between persistence diagrams and the interleaving distance between persistence modules are strongly related by the following result.

THEOREM A.20 ([CDSGO16, Theorem 5.14]). *If \mathbb{V}, \mathbb{W} are q -tame, then*

$$\mathbf{d}_\infty(\text{Dgm}(\mathbb{V}), \text{Dgm}(\mathbb{W})) = d_i(\mathbb{V}, \mathbb{W}).$$

A complementary example on stability properties of the bottleneck distance.

EXAMPLE A.21. *Let X, Y be two point clouds in \mathbb{R}^d . Consider the two filtrations $d_X = \text{dist}_X$ and $d_Y = \text{dist}_Y$. Recall (Example A.10) that f and g induce the Čech complexes of X and Y respectively. Let thus $\text{Dgm}(d_X)$ and $\text{Dgm}(d_Y)$ denote the Čech diagrams of X and Y .*

Simple computations show that

$$\begin{aligned} \|d_X - d_Y\|_\infty &= \sup_{z \in \mathbb{R}^d} \left| \min_{x \in X} \|x - z\| - \min_{y \in Y} \|y - z\| \right| \\ &= \max \left\{ \max_{y \in Y} \min_{x \in X} \|x - y\|, \max_{x \in X} \min_{y \in Y} \|x - y\| \right\} \\ &=: d_H(X, Y), \end{aligned}$$

the latter quantity being known as the Hausdorff distance between X and Y , a standard metric between point clouds (and more generally between compact metric spaces).

In this context, the stability theorem implies that

$$\mathbf{d}_\infty(\mathrm{Dgm}(d_X), \mathrm{Dgm}(d_Y)) \leq d_H(X, Y), \quad (\text{A.4})$$

that is the (bottleneck) distance between the two diagrams is controlled by the (Hausdorff) distance between the point clouds (see Figure 2.2).

This example can be extended to any finite metric spaces, in which case the stability is expressed in terms of the Gromov-Hausdorff distance or Gromov-Wasserstein distance, we refer the interested reader to [CCSG⁺09b, Theorem 3.2, Corollary 5.1].

This result has some consequences in applications. For instance, consider a (unknown) compact set $X \subset \mathbb{R}^d$, and a point cloud $P \subset X$ (for which the diagram can be computed in practice). If one can ensure that $d_H(P, X) \leq \varepsilon$ (we say that P is a ε -sample of X), then it follows that

$$\mathbf{d}_\infty(\mathrm{Dgm}(d_P), \mathrm{Dgm}(d_X)) \leq \varepsilon.$$

Under some geometric assumptions on X , not detailed here, $\mathrm{Dgm}(d_X)$ actually reflects the homology of X , and so does $\mathrm{Dgm}(d_P)$ if P is close enough to X .

APPENDIX B

ELEMENTS OF MEASURE THEORY

In the following, let Ω be a locally compact *Polish metric space* (i.e. a Polish space equipped with a distinguished Polish metric).

DEFINITION B.1 (*Radon measures*). *The space $\mathcal{M}(\Omega)$ of Radon measures supported on Ω is the space of Borel measures which give finite mass to every compact set of Ω . The vague topology on $\mathcal{M}(\Omega)$ is the coarsest topology such that the maps $\mu \mapsto \mu(f) := \int_{\Omega} f(x) d\mu(x)$ are continuous for every $f \in C_c(\Omega)$, the space of continuous functions with compact support in Ω .*

DEFINITION B.2 (*Weak topology*). *Denote by $\mathcal{M}_f(\Omega)$ the space of finite Borel measures on Ω . The weak topology on $\mathcal{M}_f(\Omega)$ is the coarsest topology such that the maps $\mu \mapsto \mu(f)$ are continuous for every $f \in C_b(\Omega)$, the space of continuous bounded functions in Ω .*

We denote by \xrightarrow{v} the vague convergence and \xrightarrow{w} the weak convergence.

DEFINITION B.3. *A set $F \subset \mathcal{M}(\Omega)$ is said to be tight if, for every $\varepsilon > 0$, there exists a compact set K with $\mu(\Omega \setminus K) \leq \varepsilon$ for every $\mu \in F$.*

The following propositions are standard results. Corresponding proofs can be found for instance in [Kal83, Section 15.7].

PROPOSITION B.4. *A set $F \subset \mathcal{M}(\Omega)$ is relatively compact for the vague topology if and only if for every compact set K included in Ω ,*

$$\sup\{\mu(K), \mu \in F\} < \infty.$$

PROPOSITION B.5 (Prokhorov's theorem). *A set $F \subset \mathcal{M}_f(\Omega)$ is relatively compact for the weak topology if and only if F is tight and $\sup_{\mu \in F} \mu(\Omega) < \infty$.*

PROPOSITION B.6. *Let μ, μ_1, μ_2, \dots be measures in $\mathcal{M}_f(\Omega)$. Then, $\mu_n \xrightarrow{w} \mu$ if and only if $\mu_n(\Omega) \rightarrow \mu(\Omega)$ and $\mu_n \xrightarrow{v} \mu$.*

PROPOSITION B.7 (Portmanteau theorem). *Let μ, μ_1, μ_2, \dots be measures in $\mathcal{M}(\Omega)$. Then, $\mu_n \xrightarrow{v} \mu$ if and only if one of the following propositions holds:*

- for all open sets $U \subset \Omega$ and all bounded closed sets $F \subset \Omega$,

$$\limsup_{n \rightarrow \infty} \mu_n(F) \leq \mu(F) \text{ and } \liminf_{n \rightarrow \infty} \mu_n(U) \geq \mu(U).$$

- for all bounded Borel sets A with $\mu(\partial A) = 0$, $\lim_{n \rightarrow \infty} \mu_n(A) = \mu(A)$.

DEFINITION B.8. *The set of point measures on Ω is the subset $\mathcal{D}(\Omega) \subset \mathcal{M}(\Omega)$ of Radon measures with discrete support and integer mass on each point, that is of the form*

$$\sum_{x \in X} n_x \delta_x$$

where $n_x \in \mathbb{N}$ and $X \subset \Omega$ is some locally finite set.

PROPOSITION B.9. *The set $\mathcal{D}(\Omega)$ is closed in $\mathcal{M}(\Omega)$ for the vague topology.*

DEFINITION B.10. *Let μ, ν be two non-negative Radon measures on a space Ω . We say that μ is absolutely continuous with respect to ν , denoted by $\mu \ll \nu$ if for all Borel $A \subset \Omega$ such that $\nu(A) = 0$, one has $\mu(A) = 0$. We say that μ and ν are singular (to each other), denoted by $\mu \perp \nu$, if there exists a Borel set $E \subset \Omega$ such that for all Borel A , $\mu(A \setminus E) = 0$ and $\nu(A \setminus E^c) = 0$.*

PROPOSITION B.11 (Radon-Nikodym derivative and Lebesgue decomposition theorem.). *Let μ, ν be two non-negative Radon measures supported on Ω . There exists a unique pair of measures (μ_1, μ_2) such that $\mu = \mu_1 + \mu_2$, $\mu_1 \ll \nu$, and $\mu_2 \perp \nu$.*

Furthermore, there exists a measurable non-negative function denoted by $\frac{d\mu_1}{d\nu}$, unique ν -ae, such that $d\mu_1 = \frac{d\mu_1}{d\nu} d\nu$. This function is called the Radon-Nikodym derivative of μ with respect to ν .

APPENDIX C

ELEMENTS OF METRIC GEOMETRY

Here, \mathcal{X} denotes a locally compact Polish metric space, equipped with a distance d . The following definitions are standard and the interested reader can find more content in [BBB⁺01] and the recent book [AKP17].

DEFINITION C.1 (Paths, lengths, geodesics). *A path between x, y in \mathcal{X} is a continuous map $\gamma : [0, 1] \rightarrow \mathcal{X}$ such that $\gamma(0) = x, \gamma(1) = y$. The length of a path γ is*

$$L(\gamma) := \inf_{0=t_1 \leq \dots \leq t_n=1} \sum_{i=1}^n d(\gamma(t_i), \gamma(t_{i+1})).$$

Let $P(x, y)$ denote the set of paths between two points x and y of \mathcal{X} .

DEFINITION C.2 (Geodesics). *The space \mathcal{X} is said to be a length space if*

$$\inf_{\gamma \in P(x, y)} L(\gamma) = d(x, y).$$

A path realizing this infimum, should it exist, is said to be a geodesic between x and y .

If for any pair $x, y \in \mathcal{X}$, there exists a geodesic between x and y , \mathcal{X} is said to be a geodesic space.

The notion of curvature in geodesic spaces is of major importance to understand the behavior of Fréchet means, among other things. For the sake of simplicity, we only expose here the notion of “non-negatively curved” Alexandrov spaces, as it covers the case of spaces of persistence diagrams. Intuitively, it means that two geodesics emanating from a same point tends

to “get closer”, as it would happen on a sphere for instance. This is measured by the notion of angle, mimicking the formula one would expect in the Euclidean setting.

DEFINITION C.3. *Let p, x, y be three points in a geodesic space \mathcal{X} . The angle $\angle_p(x, y)$ between x and y at p is defined as*

$$\cos(\angle_p(x, y)) = \frac{d(p, x)^2 + d(p, y)^2 - d(x, y)^2}{2d(p, x)d(p, y)}.$$

DEFINITION C.4. *A geodesic space \mathcal{X} is said to be a non-negatively curved Alexandrov space if for any triplet of points (p, x, y) , and for any geodesics $\gamma_x \in P(p, x)$ and $\gamma_y \in P(p, y)$, the function*

$$(s, t) \mapsto \cos(\angle_p(\gamma_x(s), \gamma_y(t)))$$

is increasing in both of its arguments.

Titre : Outils statistiques pour l'analyse topologique des données via transport optimal

Mots clés : Analyse topologique des données, Diagrammes de persistance, Transport optimal, Convergence de mesures positives, Barycentres, Algorithmes

Résumé : L'analyse topologique des données (ATD) permet d'extraire une information riche des données structurées (telles que les graphes ou les séries temporelles) présentes dans les problèmes modernes d'apprentissage. Elle va représenter cette information sous forme de descripteurs dont font partie les diagrammes de persistance, qui peuvent être décrits comme des mesures ponctuelles supportées sur un demi-plan. À défaut d'être de simples vecteurs, les diagrammes de persistance peuvent néanmoins être comparés entre eux à l'aide de métriques d'appariement partiel. La similarité entre ces métriques et les métriques usuelles du transport optimal - un autre domaine des mathématiques - est connue de longue date, mais un lien formel entre ces deux domaines restait à établir. L'objet de cette thèse est de clarifier cette connexion pour pouvoir utiliser les nom-

breux acquis du transport optimal afin de développer de nouveaux outils statistiques (théoriques et pratiques) pour manipuler les diagrammes de persistance. Dans un premier temps, nous montrons comment le transport optimal partiel avec frontière, une variante du transport optimal classique, nous fournit un formalisme qui contient les métriques usuelles de l'ATD. Nous illustrons ensuite les apports bénéfiques de cette reformulation dans différentes situations: étude théorique et algorithme pour l'estimation efficace des barycentres de diagrammes de persistance grâce au transport régularisé, caractérisation des représentations linéaires continues des diagrammes et leur apprentissage via un réseau de neurones versatile, ainsi qu'un résultat de stabilité des moyennes linéaires de diagrammes tirés aléatoirement.

Title : Statistical tools for topological descriptors via optimal transport

Keywords : Topological data analysis, Persistence diagrams, Optimal transport, Convergence of nonnegative measures, Barycenters, Algorithms

Abstract : Topological data analysis (TDA) allows one to extract rich information from structured data (such as graphs or time series) that occurs in modern machine learning problems. This information will be represented as descriptors such as persistence diagrams, which can be described as point measures supported on a half-plane. While persistence diagrams are not elements of a vector space, they can still be compared using partial matching metrics. The similarities between these metrics and those routinely used in optimal transport—another field of mathematics—are known for long, but a formal connection between these two fields is yet to come. The purpose of this thesis is to clarify this connection

and develop new theoretical and computational tools to manipulate persistence diagrams, targeting statistical applications. First, we show how optimal partial transport with boundary, a variation of classic optimal transport theory, provides a formalism that encompasses standard metrics in TDA. We then showcase the benefits of this connection in different situations: a theoretical study and the development of an algorithm to perform fast estimation of barycenters of persistence diagrams, the characterization of continuous linear representations of persistence diagrams and how to learn such representations using a neural network, and eventually a stability result in the context of linearly averaging random persistence diagrams.