

fGOT: Graph Distances Based on Filters and Optimal Transport

Hermina Petric Maretic^{1,2}, Mireille El Gheche³, Giovanni Chierchia⁴, Pascal Frossard¹

¹ EPFL, LTS4, Lausanne, Switzerland

² Amazon, London, United Kingdom

³ Sony AI, Zürich, Switzerland

⁴ Université Gustave Eiffel, LIGM (UMR 8049), ESIEE Paris, CNRS, F-93162, Noisy-le-Grand, France

hermina.petricmaretic@epfl.ch, mireille.elgheche@sony.com, giovanni.chierchia@esiee.fr, pascal.frossard@epfl.ch

Abstract

Graph comparison deals with identifying similarities and dissimilarities between graphs. A major obstacle is the unknown alignment of graphs, as well as the lack of accurate and inexpensive comparison metrics. In this work we introduce the *filter graph distance*. It is an optimal transport based distance which drives graph comparison through the probability distribution of filtered graph signals. This creates a highly flexible distance, capable of prioritising different spectral information in observed graphs, offering a wide range of choices for a comparison metric. We tackle the problem of graph alignment by computing graph permutations that minimise our new filter distances, which implicitly solves the graph comparison problem. We then propose a new approximate cost function that circumvents many computational difficulties inherent to graph comparison and permits the exploitation of fast algorithms such as mirror gradient descent, without grossly sacrificing the performance. We finally propose a novel algorithm derived from a stochastic version of mirror gradient descent, which accommodates the non-convexity of the alignment problem, offering a good trade-off between performance accuracy and speed. The experiments on graph alignment and classification show that the flexibility gained through filter graph distances can have a significant impact on performance, while the difference in speed offered by the approximation cost makes the framework applicable in practical settings.

Introduction

With the rapid development of digitization in various domains, the volume of data continuously increases. A large share of it takes the form of structured data, which is often represented by graphs that capture potentially complex information structures. The analysis of such data often requires the design of algorithms that are able to properly compute distances between different graphs. It stays however pretty challenging to compare graphs for two main reasons. First, it is usually not clear *a priori* how to align the nodes of graphs under comparison, leading to a very large number of possible solutions. Second, even if the nodes are aligned, it is not obvious how to define a meaningful metric for comparing graphs. In particular, a naive comparison of graph adjacency matrices is not an optimal measure, as

it does not capture the importance of an edge in the graph. Edges can have a very different influence on the graph structure, as well as on its local or global spectral characteristics.

In this paper, we propose the filter graph distance, a dissimilarity measure based on optimal transport, which compares two graphs through the probability distribution of data generated with graph filters. Graph filters have been largely studied in the field of Graph Signal Processing (Shuman et al. 2013; Ortega et al. 2018), and offer a high level of flexibility in modelling the relationship between data and the underlying graph.¹ Hence, their use in defining filter graph distance permits to capture a wide range of structural graph properties, including local characteristics, global structure, and any combination of spectral graph properties.

Equipped with this new distance, we derive an efficient approximation for the graph alignment problem, which is a necessary step in graph comparison. Our approximation permits to remove computational complexity bottlenecks in the alignment problem (e.g., (Petric Maretic et al. 2019)). It also enables the comparison of graphs of different size, which brings an important advantage in practice. The solution of the graph alignment problem directly leads to the computation of the minimal filter graph distance, which eventually permits to compare different graphs.

Our new formulation of the graph alignment problem leads to a nonconvex optimization framework. We solve the problem efficiently with a novel stochastic algorithm based on mirror gradient descent (MDG), which has the peculiarity of being applicable to any filter graph distance. We demonstrate the benefits of our method in alignment, clustering and classification tasks with both synthetic and real graph datasets. Filter graph distances show better performance than the standard optimal transport distances on simple graph alignment tasks, while significantly decreasing the computational cost of alignment recovery. The proposed stochastic algorithm also achieves significantly better results for community detection in structured graphs, when compared to the vanilla MGD. This suggests that our algorithm successfully addresses the inherent difficulty of graph alignment. Finally, experiments on benchmark graph classification datasets demonstrate how the benefits of our method propagate to other distance-related tasks, confirming the im-

¹The data living on a graph is also referred to as a *graph signal*.

portance of the flexibility introduced in our method.

The rest of this paper is structured as follows. The next section reviews the related work. Then we introduce the filter graph distances using the optimal transport framework and give a scalable approximation cost to the newly formulated optimal transport problem. After that, we propose a new stochastic algorithm for solving our new graph alignment problem via MGD. Finally, we assess the performance of the proposed approach in different tasks, where the benefits of the filter graph distance and the efficient stochastic algorithm is shown for synthetic and real graph datasets.

Related Work

Several works in the literature have studied the graph comparison problem. For example, graph alignment has been formulated as a quadratic assignment problem (Yan et al. 2016; Jiang et al. 2017), both as a graph edit distance (Bougleux et al. 2017) or under the constraint that the solution is a permutation matrix. This results in an NP-hard problem, and several relaxations have been proposed to find approximate solutions to this problem (Cho, Lee, and Lee 2010; Zhou and Torre 2016; Yu et al. 2018). Recently, some works have studied the graph alignment problem from the perspective of optimal transport, which has inspired several advances in machine learning (Arjovsky, Chintala, and Bottou 2017; Peyré, Cuturi et al. 2019). For example, (Gu, Hua, and Liu 2015) define a spectral distance by assigning a probability measure to the nodes via the spectrum representation of each graph and using Wasserstein distances between probability measures. This approach, however, disregards graph eigenvectors and thus does not take into account the full graph structure in the alignment problem. The authors in (Nikolentzos, Meladianos, and Vazirgiannis 2017) propose to match the graph embeddings, where the latter are represented as bags of vectors, and the Wasserstein distance is computed between them. The authors also propose a heuristic to take into account possible node labels or signals. A Gumbel-sinkhorn network inspired by optimal transport was further proposed to infer permutations from data (Mena et al. 2018; Emami and Ranka 2018). The approach consists of producing a discrete permutation from a continuous doubly-stochastic matrix obtained with the Sinkhorn operator (Sinkhorn 1964). More recently, optimal transport between stationary Markov chains has been used to find a coupling between graphs (O’Connor et al. 2021).

Building on similar ideas, the work in (Mémoli 2011) investigates the Gromov-Wasserstein distance for object matching, and the authors in (Peyré, Cuturi, and J. 2016) propose an efficient algorithm to compute the Gromov-Wasserstein distance and the barycenter of pairwise dissimilarity matrices. The algorithm uses entropic regularization and Sinkhorn projections, as proposed by (Cuturi 2013). Later, the authors in (Vayer et al. 2019) build on this work to propose a distance for graphs and signals living on these graphs. The problem is given as a combination between the Gromov-Wasserstein of graph distance matrices and the Wasserstein distance of graph signals. Then, the work in (Xu et al. 2019) proposes a method based on Gromov-Wasserstein which simultaneously learns the graph

alignment and the embeddings of graph nodes. The node embeddings are derived using optimal transport, which in turn helps in the graph matching task. The authors follow up on (Xu, Luo, and Carin 2019) by devising a scalable version of Gromov-Wasserstein distance for graph partitioning and matching. However, while the above methods solve the alignment problem using optimal transport, the simple distances between aligned graphs do not take into account the global structure of the graph, which is important for proper graph comparison.

The graph optimal transport distance introduced by (Petric Maretic et al. 2019; Petric Maretic 2021) has shown to successfully capture the global structure of graphs, representing their topology through smooth signal distributions. The work has been extended by (Maretic et al. 2020) to accommodate for comparison of graphs of different sizes using a one-to-many framework, and by (Dong and Sawin 2020) for graph sketching. However, while smooth graph signals efficiently capture global graph properties, they can lack descriptiveness for other potentially interesting properties. Several graph distances based on heat diffusion have been proposed. (Hammond, Gur, and Johnson 2013) propose a direct comparison of graph heat diffusion matrices, but stay limited to graphs of the same size. A spectral method proposed by (Tsitsulin et al. 2018) circumvents this problem by comparing heat kernel traces, but it still does not address the graph alignment problem. Closer to our work, a fast heat kernel distance based on optimal transport has recently been proposed by (Barbe et al. 2020). The authors apply the Wasserstein distance between signals only after a graph filtering step, encoding the structural information of the graph into the filtered signals. However, this distance compares graphs through available signals, whereas our method uses the representation of graphs through signal distributions and therefore does not need to rely on actual signal availability. Furthermore, the above methods are limited only to heat diffusion models, while more general distances could often be of interest in practice.

Filter Graph Alignment with Optimal Transport

Despite recent advances in the analysis of graph data, it stays pretty challenging to define a meaningful distance between graphs. Instead of comparing graphs directly, we therefore propose to look at the signal distributions governed by graph filters. Specifically, we define the filter graph distance (fGOT) as a generalisation of the graph optimal transport (GOT) distance proposed by (Petric Maretic et al. 2019), which has the ability to emphasise specific spectral properties of the graph, such as high or low frequencies, local or global graph phenomena. We model these properties through filtered graph signals, which exploit the specific graph information and eventually compare graphs through filtered signal distributions.

Preliminaries

Let $\mathcal{G} = (\mathcal{V}, \mathcal{E}, W)$ be an undirected, weighted graph with no labels and with a set of N vertices \mathcal{V} , edges \mathcal{E} and

a weighted adjacency matrix W . The combinatorial graph Laplacian is defined as $L = D - W$, where D is a diagonal matrix of node degrees. We define a signal on a graph as a vector $x \in \mathbb{R}^N$, where x_n denotes the value of a signal x on a vertex n .

We further define filtering in graph signal processing (Ortega et al. 2018) as $x_f = g(L)x$, where the filter $g(L)$ is an operator defined through the graph Laplacian matrix, and x_f represents a filtered graph signal.² A graph filter can be represented in matrix form as

$$g(L) = U\hat{G}U^T \quad (1)$$

with a diagonal matrix $\hat{G} = \text{diag}(\hat{g}(\lambda_1), \dots, \hat{g}(\lambda_N))$, where U denotes the eigenvectors and λ_i the eigenvalues of the Laplacian matrix L . Given a graph filter $g(L)$ and a Gaussian white noise signal $w \sim \mathcal{N}(0, I)$, we study random filtered graph signals given by $x_f = g(L)w$. They follow the Gaussian distribution

$$\nu^{\mathcal{G},g} := \mathcal{N}(0, g(L)I g(L)^T) = \mathcal{N}(0, g^2(L)), \quad (2)$$

where $g(L)g(L)^T = g^2(L)$ follows from the fact that every graph filter $g(L)$ is symmetric by definition in Equation 1.

Our graph distance compares graphs through their respective filtered signal distributions. The choice of filter is critical to capture fine structural properties of graphs and to drive the resulting distance. For example, a distance based on the graph heat-kernel will model differences in the spread of heat through the graphs, emphasizing global dissimilarities.

Filter Graph Distance

Given two aligned graphs \mathcal{G}_1 and \mathcal{G}_2 of the same size³ with Laplacian matrices L_1 and L_2 , we define the filter graph distance (fGOT) as the Wasserstein distance between the probabilistic distributions of their respective filtered signals, namely $\nu^{\mathcal{G}_1,g}$ and $\nu^{\mathcal{G}_2,g}$, defined earlier. In particular, the 2-Wasserstein distance measures the cumulative cost of transporting one distribution to the other in the optimal way and with respect to the Euclidean norm (Monge 1781). For normal distributions such as $\nu^{\mathcal{G}_1,g}$ and $\nu^{\mathcal{G}_2,g}$, the 2-Wasserstein distance is equal to

$$\mathcal{W}_2^2(\nu^{\mathcal{G}_1,g}, \nu^{\mathcal{G}_2,g}) = \text{Tr}(g^2(L_1)) + \text{Tr}(g^2(L_2)) - 2 \text{Tr}\left(\sqrt{g(L_1)g^2(L_2)g(L_1)}\right). \quad (3)$$

The fGOT distance compares the nature of graph filter responses, putting an emphasis on the specific properties of filtered signals. A simple example of this is shown in Figure 1. In particular, the smooth graph optimal transport (GOT) distance can be seen as a special case⁴ of fGOT, with the low pass graph filter equal to $g(L) = \sqrt{L^\dagger}$.

²The choice of the filter $g(\cdot)$ drives different spectral characteristics of the graph signals. For example, a heat kernel $g(L) = e^{-\tau L}$ models a graph based on the nature of the spread of heat through the graph, usually emphasizing global graph properties. On the other hand, a high pass filter like $g(L) = L^2$ takes more local phenomena into account, prioritising high graph frequencies.

³This assumption will be removed later.

⁴Here \dagger denotes a pseudoinverse operator.

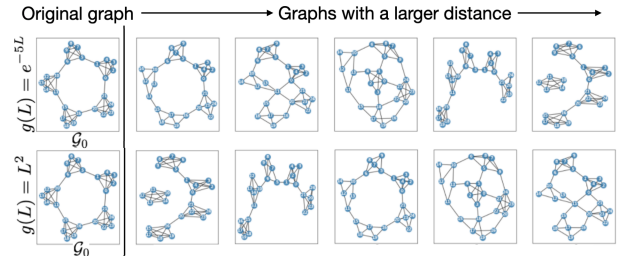


Figure 1: Graphs sorted by distance to \mathcal{G}_0 with respect to two different filter graph distances. The low-pass filter focuses on global graph properties (e.g., preservation of the ring structure, connectedness), while the high-pass filter distance considers local differences (e.g., immediate neighbourhood of each node, number of removed edges). Both can be of interest, depending on the application. For more details, please refer to (Petric Maretic 2021; Maretic et al. 2021).

Alignment Problem

It is important to note that the distance introduced in (3) requires knowledge of a correspondence between vertices of the two graphs being compared. While in some cases a consistent enumeration can be trivially chosen for all graphs, this requirement is often not realistic. Therefore, a proper alignment between graph vertices must be recovered for making graphs comparable. The alignment between two graphs \mathcal{G}_1 and \mathcal{G}_2 can be represented by a permutation matrix $P \in \mathbb{R}^{N \times N}$ that maps a vertex of \mathcal{G}_1 to a vertex of \mathcal{G}_2 . To take different alignments into account, we define a probability distribution for the permuted version of \mathcal{G}_2 as

$$\nu_P^{\mathcal{G}_2,g} = \mathcal{N}(0, g^2(PL_2P^T)). \quad (4)$$

We then aim at searching for the optimal alignment that minimises the filter graph distance between \mathcal{G}_1 and a permuted version of \mathcal{G}_2 , yielding the optimization problem defined as

$$\underset{P \in \mathcal{C}_{\text{perm}}}{\text{minimize}} \quad \mathcal{W}_2^2(\nu^{\mathcal{G}_1,g}, \nu_P^{\mathcal{G}_2,g}) \quad (5)$$

where $\mathcal{C}_{\text{perm}}$ denotes the set of permutation matrices

$$\mathcal{C}_{\text{perm}} = \left\{ P \in \mathbb{R}^{N \times N} : \begin{array}{l} (\forall i, \forall j) P_{ij} \in \{0, 1\} \\ (\forall i) \sum_j P_{ij} = 1 \\ (\forall j) \sum_i P_{ij} = 1 \end{array} \right\}. \quad (6)$$

Solving Problem (5) is computationally very challenging for several reasons, one of them being that the permutation P appears under a matrix square root. For this reason, we propose to replace the criterion in Problem (5) with an approximation through the upper bound $\widetilde{\mathcal{W}}_2^2$ of the filter graph distance, defined as

$$\widetilde{\mathcal{W}}_2^2(\nu^{\mathcal{G}_1,g}, \nu_P^{\mathcal{G}_2,g}) = \text{Tr}(g^2(L_1)) + \text{Tr}(g^2(L_2)) + - 2 \langle g(L_1)Pg(L_2), P \rangle. \quad (7)$$

The theoretical justification is provided by the next lemmas.

Lemma 1. *Let $P \in \mathcal{C}_{\text{perm}}$ be a permutation matrix, and let $L \in \mathbb{R}^{N \times N}$ be a graph Laplacian matrix. Then $g(PLP^T) = Pg(L)P^T$ for any graph filter $g(\cdot)$ defined as in (1).*

Proof. Please refer to (Maretic et al. 2021). \square

Lemma 2. Let \mathcal{G}_1 and \mathcal{G}_2 be two graphs with their respective Laplacian matrices $L_1 \in \mathbb{R}^{N \times N}$ and $L_2 \in \mathbb{R}^{N \times N}$. Let the fGOT distance \mathcal{W}_2^2 be defined as in Equation 3, and its approximation $\widetilde{\mathcal{W}}_2^2$ as in Equation 7. Then, for any graph filter $g(\cdot)$ defined as in (1), and $P \in \mathcal{C}_{\text{perm}}$:

$$\mathcal{W}_2^2(\nu^{\mathcal{G}_1, g}, \nu_{P^T}^{\mathcal{G}_2, g}) \leq \widetilde{\mathcal{W}}_2^2(\nu^{\mathcal{G}_1, g}, \nu_P^{\mathcal{G}_2, g}). \quad (8)$$

Proof. Please refer to (Maretic et al. 2021). \square

Lemma 3. Let \mathcal{G}_1 and \mathcal{G}_2 be two isomorphic graphs with their respective Laplacian matrices $L_1 \in \mathbb{R}^{N \times N}$ and $L_2 \in \mathbb{R}^{N \times N}$. Let the fGOT distance \mathcal{W}_2^2 be defined as in Equation 3, and its approximation $\widetilde{\mathcal{W}}_2^2$ as in Equation 7. Then, for any graph filter $g(\cdot)$ defined as in (1):

$$\min_{P \in \mathcal{C}_{\text{perm}}} \mathcal{W}_2^2(\nu^{\mathcal{G}_1, g}, \nu_{P^T}^{\mathcal{G}_2, g}) = \min_{P \in \mathcal{C}_{\text{perm}}} \widetilde{\mathcal{W}}_2^2(\nu^{\mathcal{G}_1, g}, \nu_P^{\mathcal{G}_2, g}) = 0. \quad (9)$$

Proof. Please refer to (Maretic et al. 2021). \square

In the next section, we propose an efficient algorithm to find an assignment matrix P that minimizes the surrogate cost function in (7), whose purpose is to approximately solve Problem (5). In addition to being computationally efficient, the proposed approach can advantageously handle the comparison of graphs with different sizes, because the cost function in Problem (5) and its surrogate version in Equation (7) still hold when the assignment matrix P is rectangular.

FGOT Algorithm

Relaxed Alignment Problem

Equipped with the surrogate fGOT distance in (7) to compare unaligned graphs, we now introduce the proposed formulation for approximate graph alignment. The last roadblock in solving the original graph alignment problem of Eq. (5) is the requirement that the sought solution must be a binary assignment matrix, since this criterion leads to a discrete optimisation problem with a factorial number of feasible solutions. We propose to circumvent this issue by relaxing the constraint through an implicit reformulation. Given two graphs \mathcal{G}_1 and \mathcal{G}_2 of possibly *different* size, we represent a soft assignment between the graph nodes through a matrix of the following convex set

$$\mathcal{C}_{\text{fuzzy}} = \left\{ P \in \mathbb{R}^{|V_1| \times |V_2|} : \begin{array}{l} (\forall i, \forall j) P_{ij} \geq 0 \\ (\forall i) \sum_j P_{ij} = \frac{1}{|V_1|} \\ (\forall j) \sum_i P_{ij} = \frac{1}{|V_2|} \end{array} \right\}. \quad (10)$$

Note that $P \in \mathcal{C}_{\text{perm}}$ implies $P \in \mathcal{C}_{\text{fuzzy}}$ up to a rescaling factor of $\sqrt{|V_1||V_2|}$. Using the surrogate fGOT distance in (7), we can now express the approximate distance \mathcal{D} between graphs \mathcal{G}_1 and \mathcal{G}_2 as the solution to the following graph alignment problem

$$\mathcal{D}(\mathcal{G}_1, \mathcal{G}_2) = \min_{P \in \mathcal{C}_{\text{fuzzy}}} \widetilde{\mathcal{W}}_2^2(\nu^{\mathcal{G}_1, g}, \nu_P^{\mathcal{G}_2, g}) \quad (11)$$

The standard approach to solve the above problem is by using the projected mirror gradient descent (MGD), where the projection onto $\mathcal{C}_{\text{fuzzy}}$ is computed according to the Kullback-Leibler (KL) metric (Benamou et al. 2015). Adding the entropic regularisation $\epsilon H(P)$ to the problem yields the following iterative algorithm

$$P_{t+1} = \mathcal{P}_{\mathcal{C}_{\text{fuzzy}}}^{\text{KL}}(P_t \odot e^{-\alpha q_t}), \quad (12)$$

where $q_t = \nabla \widetilde{\mathcal{W}}_2^2(\nu^{\mathcal{G}_1, g}, \nu_{P_t}^{\mathcal{G}_2, g}) - \epsilon \nabla H(P_t)$, the symbol \odot is the Haddamard (pointwise) product between matrices, ∇ represents the gradient operator, α is the step size and ϵ is the entropic regularisation parameter. The KL projection $\mathcal{P}_{\mathcal{C}_{\text{fuzzy}}}^{\text{KL}}$ can be computed through the Sinkhorn operator \mathcal{S}_τ as (Sinkhorn 1964), (Cuturi 2013)

$$\mathcal{P}_{\mathcal{C}_{\text{fuzzy}}}^{\text{KL}}(P) = \mathcal{S}_\tau(-\tau \log P), \quad (13)$$

where $\tau > 0$ is a small constant.

This is the standard approach commonly found in the literature. However, Problem (11) is non-convex and thus very susceptible to converge towards locally optimal solutions. To address this issue, we propose a novel stochastic version of MGD.

Proposed Algorithm

We reformulate Problem (11) with an implicit constraint, taking into account the KL projection appearing in the non-stochastic version of MGD algorithm. For the sake of clarity, we denote the KL projection to $\mathcal{C}_{\text{fuzzy}}$ with $\mathcal{B}(P) = \mathcal{P}_{\mathcal{C}_{\text{fuzzy}}}^{\text{KL}}(P)$, and the optimization problem becomes

$$\underset{P \in \mathbb{R}^{|V_1| \times |V_2|}}{\text{minimize}} \quad \widetilde{\mathcal{W}}_2^2(\nu^{\mathcal{G}_1, g}, \nu_{\mathcal{B}(P)}^{\mathcal{G}_2, g}). \quad (14)$$

Note that the KL projection is computed as a sequence of softmax operations, so the above cost function is constraint-free and differentiable (Luise et al. 2018), albeit nonconvex. To deal with nonconvexity, we propose to optimise the expectation of $\widetilde{\mathcal{W}}_2^2(\nu^{\mathcal{G}_1, g}, \nu_{\mathcal{B}(P)}^{\mathcal{G}_2, g})$ w.r.t. the parameters θ of some distribution p_θ , that is

$$\underset{\theta}{\text{minimize}} \quad \underbrace{\mathbb{E}_{P \sim p_\theta} \{ \widetilde{\mathcal{W}}_2^2(\nu^{\mathcal{G}_1, g}, \nu_{\mathcal{B}(P)}^{\mathcal{G}_2, g}) \}}_{U(\theta)}. \quad (15)$$

The reformulated problem aims at shaping the distribution p_θ so as to put all its mass on a minimizer P^* of $\widetilde{\mathcal{W}}_2^2$, thus integrating the use of Bayesian exploration in the optimization process. This approach is commonly used in many areas of stochastic search, such as evolution strategies (Hansen and Ostermeier 2001; Wierstra et al. 2008; Salimans et al. 2017).

Since the cost function in Problem (15) is differentiable, a critical point can be found with the gradient method, whose iterations can then be written as

$$\theta_{t+1} = \underset{\theta}{\text{argmin}} \quad \theta^\top \nabla U(\theta_t) + \frac{1}{2\alpha_t} \|\theta - \theta_t\|^2. \quad (16)$$

Replacing the Euclidean distance with a Bregman divergence \mathbb{D} leads to the algorithm known as mirror descent

$$\theta_{t+1} = \underset{\theta}{\text{argmin}} \quad \theta^\top \nabla U(\theta_t) + \frac{1}{\alpha_t} \mathbb{D}(\theta, \theta_t). \quad (17)$$

Algorithm 1: Approximate solution to Problem (11)

- 1: **Input:** Graphs \mathcal{G}_1 and \mathcal{G}_2
- 2: **Input:** Sampling $S \in \mathbb{N}$, step size $\alpha > 0$ and $\tau > 0$
- 3: **Input:** Setting matrix η_0 (constant) and σ_0 (random)
- 4: **for** $t = 0, 1, \dots$ **do**
- 5: Draw the samples $\epsilon_t^{(1)}, \dots, \epsilon_t^{(N)}$ from $\mathcal{N}(0, I)$.
- 6: Estimate the gradient $(\mathbf{g}_t^{(\eta)}, \mathbf{g}_t^{(\sigma)})$:

$$\mathbf{g}_t^{(\eta)} \approx \frac{1}{N} \sum_{n=1}^N \nabla \widetilde{\mathcal{W}}_2^2(\nu^{\mathcal{G}_1, g}, \nu_{\mathcal{B}(P_n)}^{\mathcal{G}_2, g}) \Big|_{P_n = \eta_t + \sigma_t \circ \epsilon_t^{(n)}}$$

$$\mathbf{g}_t^{(\sigma)} \approx \frac{1}{N} \sum_{n=1}^N \left(\epsilon_t^{(n)} \circ \nabla \widetilde{\mathcal{W}}_2^2(\nu^{\mathcal{G}_1, g}, \nu_{\mathcal{B}(P_n)}^{\mathcal{G}_2, g}) \Big|_{P_n = \eta_t + \sigma_t \circ \epsilon_t^{(n)}} \right)$$

- 7: Update η_t and σ_t using $(\mathbf{g}_t^{(\eta)}, \mathbf{g}_t^{(\sigma)})$:

$$\eta_{t+1} = \eta_t - \alpha_t \sigma_t^2 \circ \mathbf{g}_t^{(\eta)}, \quad \sigma_{t+1} = \sqrt{\sigma_t^2 + \mathbf{d}_t^2} - \mathbf{d}_t$$

$$\text{with } \mathbf{d}_t = \frac{1}{2} \alpha_t \sigma_t^2 \circ \mathbf{g}_t^{(\sigma)}.$$

- 8: **end for**

- 9: **Output:** $P = \mathcal{B}(\eta_*)$
-

For exponential-family distributions, the Kullback-Leibler (KL) divergence is a Bregman divergence (Raskutti and Mukherjee 2015), which is better indicated to optimize the parameters of a probability distribution (Amari 1998):

$$\mathbb{D}(\theta, \theta_t) = \mathbb{D}_{\text{KL}}(p_\theta \parallel p_{\theta_t}) = \mathbb{E}_{P \sim p_\theta} \left\{ \log \frac{p_\theta(P)}{p_{\theta_t}(P)} \right\}. \quad (18)$$

Among the possible choices for the distribution p_θ , one of the most appropriate for our scenario is the diagonal Gaussian distribution $\mathcal{N}(\eta, \text{diag}(\sigma^2))$ with $\theta = (\eta, \sigma)$, because the KL divergence is given by

$$\mathbb{D}_{\text{KL}}(p_{\eta, \sigma} \parallel p_{\eta_t, \sigma_t}) = \frac{1}{2} \sum_{d=1}^D \left(\frac{\sigma_d^2}{\sigma_{t,d}^2} + \frac{(\eta_{t,d} - \eta_d)^2}{\sigma_{t,d}^2} - 1 + \log \frac{\sigma_{t,d}^2}{\sigma_d^2} \right). \quad (19)$$

Then, the mirror-descent update boils down to

$$\eta_{t+1,d} = \eta_{t,d} - \alpha_t \sigma_{t,d}^2 \frac{\partial U(\theta_t)}{\partial \eta_d}$$

$$\sigma_{t+1,d} = \sqrt{\sigma_{t,d}^2 + \left(\frac{\alpha_t \sigma_{t,d}^2}{2} \frac{\partial U(\theta_t)}{\partial \sigma_d} \right)^2} - \frac{\alpha_t \sigma_{t,d}^2}{2} \frac{\partial U(\theta_t)}{\partial \sigma_d}. \quad (20)$$

The stochastic MGD algorithm for fGOT is summarised in Algorithm 1. Note that its computational complexity boils down to that of the matrix multiplications for each iteration of the algorithm. These can be computed very efficiently in the presence of sparse graphs, which is a common case in practice. We implemented Algorithm 1 in PyTorch using the AMSGrad method (Reddi, Kale, and Kumar 2018). Our full implementation, using some functions from (Flamary et al. 2021), is available at <https://github.com/Hermina/fGOT>.




Figure 2: Performance comparison of different algorithms and filters on a task of random graph alignment. The performance is shown in terms of the Frobenius distance between aligned graph Laplacian matrices across different graph sizes.




Figure 3: Speed comparison (natural log of execution time in *seconds*) of different algorithms and filters on a task of random graph alignment across different graph sizes.

Experimental Results

We now illustrate the behaviour of fGOT in different experimental challenges. We consider three sets of experiments. The first one deals with random unstructured graphs, which we represent by random Erdos-Renyi graphs. We demonstrate the benefits of using high-pass filters to define a graph distance, and show the speed advancement compared to the GOT algorithm presented in (Petric Maretic et al. 2019). In the second set of experiments, we consider highly structured graphs, represented by stochastic block models. We tackle the problem of matching communities and show the benefits of low-pass filters. We also highlight the superior performance of our novel stochastic algorithm. Finally, we evaluate our method on a benchmark graph classification task, and compare to several state-of-the-art methods for graph alignment. The hyperparameters for all experiments were set empirically (for details, please consult (Maretic et al. 2021)).

Alignment of Unstructured Graphs

In our first experiment, we tackle the problem of alignment of unstructured random graphs across different graph sizes. We compare the alignment quality, as well as the execution time of different algorithms and filters. In particular, for each predefined graph size between 10 and 100, we perform 50 repetitions of aligning two random Erdos-Renyi graphs using the following algorithms: Gromov-Wasserstein (GW) (Peyré, Cuturi, and J. 2016), with graphs compared based on their shortest path matrices, as proposed in (Vayer et al. 2019); three algorithms solving for the original GOT cost: the original stochastic GOT algorithm, as proposed in (Petric Maretic et al. 2019), and our fGOT and stochastic

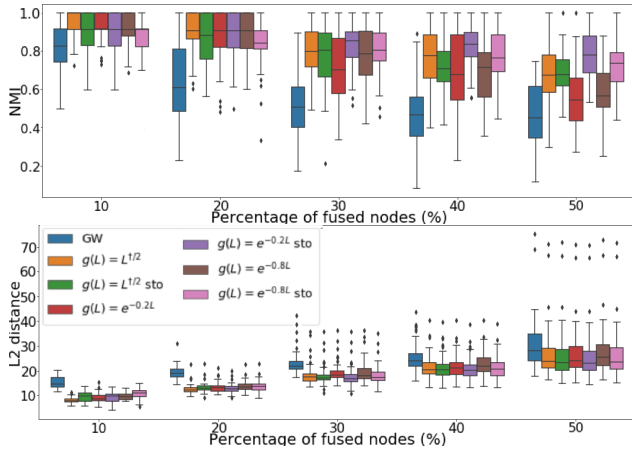


Figure 4: Alignment and community detection performance for distorted stochastic block model graphs as a function of the percentage of fused nodes (Experiment 1). The first plot shows the community detection performance in terms of NMI (closer to 1 is better), while the second one shows the Frobenius distance between aligned graph Laplacian matrices (closer to 0 is better).

fGOT algorithms, using the corresponding low pass filter $g(L) = \sqrt{L^\dagger}$; and two algorithms solving for a high-pass filter cost: fGOT and stochastic fGOT with $g(L) = L^2$.

The task of unstructured graph alignment is better suited for distances defined by a high pass filter, as shown in Figure 2. This is because high-pass filters prioritise local graph properties, which are crucial in the case of unstructured graphs due to the lack of global structure. This reinforces the importance of the right filter choice for each problem, and emphasizes the benefits of the flexibility offered by fGOT.

Furthermore, we demonstrate the execution speed of each algorithm across different graph sizes. Compared to the original GOT formulation presented in (Petric Maretic et al. 2019), we can observe a significant speed improvement in both of our fGOT algorithms. In particular, using the corresponding filter $g(L) = \sqrt{L^\dagger}$, the MGD algorithm renders the computation of the original GOT distance comparable to the GW distance speed, circumventing the largest drawback of the GOT framework. The stochastic MGD algorithm also offers a great speed improvement, while successfully accounting for the nonconvexity of the alignment problem. This allows the algorithm to capitalise on both speed and accuracy (in terms of lower l_2 distance), offering better performance than its non-stochastic counterpart at a faster rate than GOT.

In summary, our two proposed algorithms offer a choice between maximising accuracy or speed, while both provide a very competitive trade-off between the two. Furthermore, our fGOT framework highlights the benefits of using filters in graph distances and offers high flexibility in choosing the appropriate filter for the problem at hand. For instance, unstructured graph alignment calls for high pass filters, while low-pass filters are more suitable for structured graphs, as we will see in the next section.

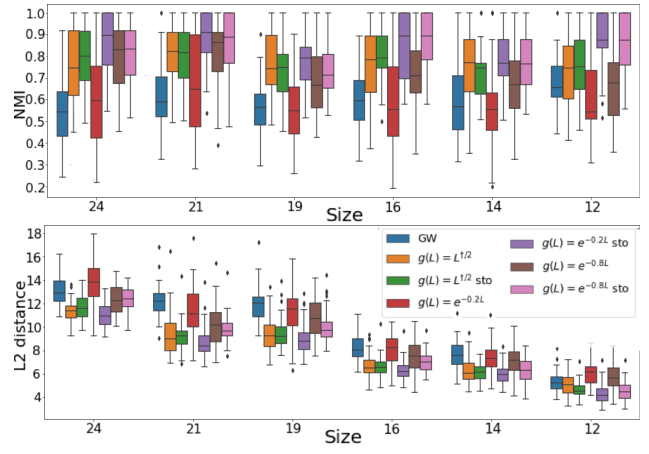


Figure 5: Alignment and community detection performance for random instances of stochastic block model graphs as a function of the graph size (Experiment 2). The first plot shows the community detection performance in terms of NMI (closer to 1 is better), while the second one shows the Frobenius distance between aligned graph Laplacian matrices (closer to 0 is better).

Community Detection in Structured Graphs

We test the performance of fGOT on a graph alignment and community detection task in structured graphs. We consider two experimental settings: the comparison of a graph with a noisy version of the same graph, and its comparison with a random structured graph. We show the results obtained by GW (as described in the previous section) and by the MGD and the stochastic MGD version of fGOT for three different graph filter models. Namely $g(L) = \sqrt{L^\dagger}$, which corresponds to the approximated version of the original GOT distance for smooth signals (Petric Maretic et al. 2019), and two heat kernel distances: $g(L) = e^{-0.2L}$ and $g(L) = e^{-0.8L}$.

To evaluate the alignment recovery, we resort to two measures: Normalised mutual information (NMI) of community alignment, and the difference between the aligned graphs in terms of the l_2 norm. We estimate the community alignment of these structured graphs by performing spectral clustering in both (aligned) graphs, separating their nodes into 4 communities. The two measures can be considered as complementary, with NMI capturing the recovery of global structure, and l_2 considering local differences independently.

In Experiment 1, we generate a stochastic block model graph \mathcal{G}_2 with 24 nodes and 4 communities. The graph \mathcal{G}_1 is constructed as a noisy version of \mathcal{G}_2 by randomly collapsing edges until a target percentage of nodes is fused. Graph \mathcal{G}_1 is then further randomly permuted, in order to change the order of its nodes. Each experiment has been repeated 50 times with a different \mathcal{G}_2 . The results are given in Figure 4.

In Experiment 2, both graphs are generated randomly as instances of stochastic block models with 4 communities. For each \mathcal{G}_2 with size 24, we generate six graphs \mathcal{G}_1 with a different number of edges and vertices (between 12 and 24). While we keep the number of communities equal for a

| | BZR | MUTAG | ENZYMES | PROTEIN | AIDS |
|---|---------------------|---------------------|--------------------|---------------------|---------------------|
| $g_1(L) = \sqrt{L^\dagger}$ | 77.7 ± 7.17 | 85.56 ± 5.98 | 27.89 ± 9.01 | 64.56 ± 9.56 | 99.0 ± 1.95 |
| $g_2(L) = L^2$ | 82.07 ± 6.34 | 84.56 ± 6.69 | 26.0 ± 8.71 | 61.0 ± 9.32 | 96.89 ± 3.44 |
| $g_3(L) = e^{-0.8L}$ | 79.43 ± 7.17 | 86.78 ± 5.87 | 25.56 ± 7.06 | 62.44 ± 9.92 | 98.78 ± 2.02 |
| $g_4(L) = \sqrt{L^\dagger} + L^2$ | 78.05 ± 11.36 | 84.44 ± 6.8 | 24.33 ± 9.89 | 66.22 ± 9.69 | 99.44 ± 1.24 |
| $g_5(L) = \sqrt{L^\dagger} + e^{-0.8L}$ | 78.39 ± 6.53 | 86.44 ± 5.37 | 29.0 ± 8.74 | 64.67 ± 9.13 | 99.22 ± 1.41 |
| $g_6(L) = L^2 + e^{-0.8L}$ | 77.01 ± 10.95 | 83.22 ± 7.58 | 26.67 ± 7.2 | 63.56 ± 8.47 | 99.44 ± 1.51 |
| GW | 79.22 ± 7.68 | 84.11 ± 6.98 | 24.67 ± 5.94 | 61.22 ± 9.33 | 95.0 ± 4.28 |
| GA | 72.11 ± 9.69 | 70.33 ± 8.09 | 17.44 ± 6.17 | 47.89 ± 11 | 82.56 ± 8.29 |
| IPFP | 69 ± 8.54 | 77.89 ± 6.92 | 21.22 ± 6.34 | 49.67 ± 7.5 | 84.22 ± 7.68 |
| RRWM | 66.11 ± 12.66 | 64.44 ± 10.55 | 20.44 ± 8.292 | 47.22 ± 11.58 | 81.33 ± 8.824 |

Table 1: 1-NN classification accuracy

meaningful community detection task, there is no other direct connection between \mathcal{G}_1 and \mathcal{G}_2 . Each experiment was repeated 50 times with a different \mathcal{G}_2 . The results are in Figure 5.

Both experiments demonstrate the superior performance of our stochastic algorithm compared to the non-stochastic MGD algorithm, especially in the case of heat kernel models. The improvements are especially visible in the community detection tasks, as the settings become more challenging. Furthermore, the stochastic version of the heat kernel distances outperforms both the approximated GOT distance, as well as the GW in both experiments. This shows that, even when observing only smooth filters, the flexibility offered by considering different filter models can affect our results substantially and bring considerable benefits. Finally, the results in terms of the l_2 distance are consistent with the NMI, suggesting that the stochastic algorithm improves the alignment significantly both in terms of global structure recovery and more local alignment properties.

Graph Classification

Finally, we test the effect of filter graph distances in a graph classification task. We explore five different benchmark datasets: BZR dataset (average number of nodes $\bar{N} = 35.75$) (Sutherland, O’Brien, and Weaver 2003), MUTAG ($\bar{N} = 17.93$) (Debnath et al. 1991), ENZYMES ($\bar{N} = 32.63$), PROTEIN ($\bar{N} = 39.06$) (Borgwardt and Kriegel 2005) and AIDS ($\bar{N} = 15.69$). From each dataset, we sample 100 graphs, use fGOT to align them and to compute normalised pairwise distances, and then perform a simple non-parametric 1-NN classification. We repeat this experiment 30 times and report the result in terms of the mean and standard deviation of the classification accuracy. We compare several variants of fGOT, emphasising the flexibility of the method and the benefits this flexibility can bring when comparing datasets of different nature. The filters in comparison are as follows: $g_1(L) = \sqrt{L^\dagger}$, $g_2(L) = L^2$, $g_3(L) = e^{-0.8L}$, $g_4(L) = \sqrt{L^\dagger} + L^2$, $g_5(L) = \sqrt{L^\dagger} + e^{-0.8L}$, $g_6(L) = L^2 + e^{-0.8L}$. These variants of fGOT are further compared to several state-of-the-art graph alignment algorithms: GW (Peyré, Cuturi, and J. 2016; Vayer et al. 2019), GA (Gold and Rangarajan 1996), IPFP (Leordeanu, Hebert, and Sukthankar 2009) and RRWM (Cho, Lee, and Lee 2010). As shown in Table 1, a high-pass filter (eg.

$g_2(L)$) will have a higher classification accuracy on a dataset with local changes, such as BZR. Conversely, a low-pass filter (eg. $g_1(L)$, $g_3(L)$) will perform better in a dataset with strong global differences, such as AIDS. This reinforces the value of the flexibility that fGOT offers, and allows us to choose the right distance criteria for our data. Even in the more challenging task of 6-class classification (ENZYMES), the fGOT distances seem to be more relevant, even though we note that 1-NN classification is probably too simplistic to solve such difficult tasks. Finally, composite filters ($g_4(L)$, $g_5(L)$, $g_6(L)$), covering both low and high frequencies, perform competitively on all observed datasets, suggesting their efficiency in retaining different captured effects. This makes them general enough to use on data for which we might not have a clear filter preference, or when we are unaware of the nature and properties of our graphs.

Conclusion

We proposed fGOT, a flexible method for graph comparison based on optimal transport, which uses filter models to encode specific graph properties. We exploit the representation of graphs through the distribution of filtered signals, which permits an efficient comparison through the Wasserstein distance of these distributions. In order to provide a more scalable algorithm, we formulate an approximation to the generic filter optimal transport cost, circumventing the computation of its most expensive parts. We propose an efficient stochastic algorithm based on Bayesian exploration, which adapts mirror gradient descent to this challenging non-convex problem. We show the performance of our method in the context of graph alignment, community detection and classification. Experimental results show that both our proposed algorithms offer a superior performance in terms of speed, when compared to the original GOT algorithm. Furthermore, our novel stochastic algorithm offers superior performance to its non-stochastic counterpart in terms of accuracy and alignment quality, reaching a very good trade-off between speed and accuracy. Finally, experiments on both unstructured and structured graphs show that the filter graph distance brings valuable flexibility, with problem-adjusted filters leading to better performance than classical metrics in different simple tasks.

References

- Amari, S.-I. 1998. Natural Gradient Works Efficiently in Learning. *Neural Comput.*, 10(2): 251–276.
- Arjovsky, M.; Chintala, S.; and Bottou, L. 2017. Wasserstein generative adversarial networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, 214–223.
- Barbe, A.; Sebban, M.; Gonçalves, P.; Borgnat, P.; and Grignonval, R. 2020. Graph Diffusion Wasserstein Distances. In *European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases*.
- Benamou, J.-D.; Carlier, G.; Cuturi, M.; Nenna, L.; and Peyré, G. 2015. Iterative Bregman Projections for Regularized Transportation Problems. *SIAM Journal on Scientific Computing*, 37(2): A1111–A1138.
- Borgwardt, K. M.; and Kriegel, H.-P. 2005. Shortest-path kernels on graphs. In *Fifth IEEE international conference on data mining (ICDM'05)*, 8–pp. IEEE.
- Bougleux, S.; Brun, L.; Carletti, V.; Foggia, P.; Gaüzère, B.; and Vento, M. 2017. Graph edit distance as a quadratic assignment problem. *Pattern Recognition Letters*, 87: 38–46.
- Cho, M.; Lee, J.; and Lee, K. M. 2010. Reweighted random walks for graph matching. In *European conference on Computer vision*, 492–505. Springer.
- Cuturi, M. 2013. Sinkhorn Distances: Lightspeed Computation of Optimal Transport. In Burges, C. J. C.; Bottou, L.; Welling, M.; Ghahramani, Z.; and Weinberger, K. Q., eds., *Advances in Neural Information Processing Systems*, 2292–2300. Curran Associates, Inc.
- Debnath, A. K.; Lopez de Compadre, R. L.; Debnath, G.; Shusterman, A. J.; and Hansch, C. 1991. Structure-activity relationship of mutagenic aromatic and heteroaromatic nitro compounds. correlation with molecular orbital energies and hydrophobicity. *Journal of medicinal chemistry*, 34(2): 786–797.
- Dong, Y.; and Sawin, W. 2020. COPT: Coordinated Optimal Transport for Graph Sketching. *arXiv preprint arXiv:2003.03892*.
- Emami, P.; and Ranka, S. 2018. Learning Permutations with Sinkhorn Policy Gradient. *Preprint arXiv:1805.07010*.
- Flamary, R.; Courty, N.; Gramfort, A.; Alaya, M. Z.; Boisbunon, A.; Chambon, S.; Chapel, L.; Corenflos, A.; Fatras, K.; Fournier, N.; Gautheron, L.; Gayraud, N. T.; Janati, H.; Rakotomamonjy, A.; Redko, I.; Rolet, A.; Schutz, A.; Seguy, V.; Sutherland, D. J.; Tavenard, R.; Tong, A.; and Vayer, T. 2021. POT: Python Optimal Transport. *Journal of Machine Learning Research*, 22(78): 1–8.
- Gold, S.; and Rangarajan, A. 1996. A graduated assignment algorithm for graph matching. *IEEE Transactions on pattern analysis and machine intelligence*, 18(4): 377–388.
- Gu, J.; Hua, B.; and Liu, S. 2015. Spectral distances on graphs. *Discrete Applied Mathematics*, 190–191: 56–74.
- Hammond, D. K.; Gur, Y.; and Johnson, C. R. 2013. Graph diffusion distance: A difference measure for weighted graphs based on the graph Laplacian exponential kernel. In *2013 IEEE Global Conference on Signal and Information Processing*, 419–422. IEEE.
- Hansen, N.; and Ostermeier, A. 2001. Completely De-randomized Self-Adaptation in Evolution Strategies. *Evol. Comput.*, 9(2): 159–195.
- Jiang, B.; Tang, J.; Ding, C.; Gong, Y.; and Luo, B. 2017. Graph Matching via Multiplicative Update Algorithm. In Guyon, I.; Luxburg, U. V.; Bengio, S.; Wallach, H.; Fergus, R.; Vishwanathan, S.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, 3187–3195. Curran Associates, Inc.
- Leordeanu, M.; Hebert, M.; and Sukthankar, R. 2009. An Integer Projected Fixed Point Method for Graph Matching and MAP Inference. In Bengio, Y.; Schuurmans, D.; Lafferty, J. D.; Williams, C. K. I.; and Culotta, A., eds., *Advances in Neural Information Processing Systems 22*, 1114–1122. Curran Associates, Inc.
- Luise, G.; Rudi, A.; Pontil, M.; and Ciliberto, C. 2018. Differential Properties of Sinkhorn Approximation for Learning with Wasserstein Distance. In Bengio, S.; Wallach, H.; Larochelle, H.; Grauman, K.; Cesa-Bianchi, N.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, 5859–5870. Curran Associates, Inc.
- Maretic, H. P.; Gheche, M. E.; Chierchia, G.; and Frossard, P. 2021. FGOT: Graph Distances based on Filters and Optimal Transport. *arXiv:2109.04442*.
- Maretic, H. P.; Gheche, M. E.; Minder, M.; Chierchia, G.; and Frossard, P. 2020. Wasserstein-based Graph Alignment. *arXiv preprint arXiv:2003.06048*.
- Mémoli, F. 2011. Gromov–Wasserstein distances and the metric approach to object matching. *Foundations of computational mathematics*, 11(4): 417–487.
- Mena, G.; Belanger, D.; Linderman, S.; and Snoek, J. 2018. Learning Latent Permutations with Gumbel-Sinkhorn Networks. In *International Conference on Learning Representations*.
- Monge, M. 1781. *Mémoire sur la théorie des déblais et des remblais*. De l’Imprimerie Royale.
- Nikolentzos, G.; Meladianos, P.; and Vazirgiannis, M. 2017. Matching node embeddings for graph similarity. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- O’Connor, K.; Yi, B.; McGoff, K.; and Nobel, A. B. 2021. Graph Optimal Transport with Transition Couplings of Random Walks. *arXiv:2106.07106 [cs, stat]*. ArXiv: 2106.07106.
- Ortega, A.; Frossard, P.; Kovačević, J.; Moura, J. M.; and Vandergheynst, P. 2018. Graph signal processing: Overview, challenges, and applications. *Proceedings of the IEEE*, 106(5): 808–828.
- Petric Maretic, H. 2021. *Representing graphs through data with learning and optimal transport*. Ph.D. thesis, École Polytechnique Fédérale de Lausanne (EPFL).
- Petric Maretic, H.; El Gheche, M.; Chierchia, G.; and Frossard, P. 2019. GOT: An Optimal Transport framework for Graph comparison. In Wallach, H.; Larochelle, H.; Beygelzimer, A.; d’Alché Buc, F.; Fox, E.; and Garnett, R.,

- eds., *Advances in Neural Information Processing Systems* 32, 13876–13887. Curran Associates, Inc.
- Peyré, G.; Cuturi, M.; and J., S. 2016. Gromov-Wasserstein Averaging of Kernel and Distance Matrices. In Balcan, M. F.; and Weinberger, K. Q., eds., *International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, 2664–2672. New York, New York, USA.
- Peyré, G.; Cuturi, M.; et al. 2019. Computational Optimal Transport: With Applications to Data Science. *Foundations and Trends® in Machine Learning*, 11(5-6): 355–607.
- Raskutti, G.; and Mukherjee, S. 2015. The information geometry of mirror descent. *IEEE Transactions on Information Theory*, 61(3): 1451–1457.
- Reddi, S. J.; Kale, S.; and Kumar, S. 2018. On the Convergence of Adam and Beyond. In *International Conference on Learning Representations*.
- Salimans, T.; Ho, J.; Chen, X.; Sidor, S.; and Sutskever, I. 2017. Evolution strategies as a scalable alternative to reinforcement learning. *arXiv preprint arXiv:1703.03864*.
- Shuman, D. I.; Narang, S. K.; Frossard, P.; Ortega, A.; and Vandergheynst, P. 2013. The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains. *IEEE Signal Processing Magazine*, 30(3): 83–98.
- Sinkhorn, R. 1964. A relationship between arbitrary positive matrices and doubly stochastic matrices. *The Annals of Mathematical Statistics*, 35(2): 876–879.
- Sutherland, J. J.; O’Brien, L. A.; and Weaver, D. F. 2003. Spline-fitting with a genetic algorithm: A method for developing classification structure- activity relationships. *Journal of chemical information and computer sciences*, 43(6): 1906–1915.
- Tsitsulin, A.; Mottin, D.; Karras, P.; Bronstein, A.; and Müller, E. 2018. Netlsd: hearing the shape of a graph. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2347–2356. ACM.
- Vayer, T.; Chapel, L.; Flamary, R.; Tavenard, R.; and Courty, N. 2019. Optimal Transport for structured data with application on graphs. In *ICML 2019-36th International Conference on Machine Learning*, 1–16.
- Wierstra, D.; Schaul, T.; Peters, J.; and Schmidhuber, J. 2008. Natural Evolution Strategies. In *2008 IEEE Congress on Evolutionary Computation (IEEE World Congress on Computational Intelligence)*, 3381–3387.
- Xu, H.; Luo, D.; and Carin, L. 2019. Scalable Gromov-Wasserstein Learning for Graph Partitioning and Matching. In Wallach, H.; Larochelle, H.; Beygelzimer, A.; d’Alché Buc, F.; Fox, E.; and Garnett, R., eds., *Advances in Neural Information Processing Systems* 32, 3046–3056. Curran Associates, Inc.
- Xu, H.; Luo, D.; Zha, H.; and Duke, L. C. 2019. Gromov-Wasserstein Learning for Graph Matching and Node Embedding. In *International Conference on Machine Learning*, 6932–6941.
- Yan, J.; Yin, X.; Lin, W.; Deng, C.; Zha, H.; and Yang, X. 2016. A short survey of recent advances in graph matching. In *International Conference on Multimedia Retrieval*, 167–174. New York, NY, USA: ACM.
- Yu, T.; Yan, J.; Wang, Y.; Liu, W.; and Li, B. 2018. Generalizing Graph Matching beyond Quadratic Assignment Model. In Bengio, S.; Wallach, H.; Larochelle, H.; Grauman, K.; Cesa-Bianchi, N.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, 853–863. Curran Associates, Inc.
- Zhou, F.; and Torre, F. D. 2016. Factorized Graph Matching. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(9): 1774–1789.