

A Family of Tractable Graph Metrics

José Bento¹ and Stratis Ioannidis^{2*}

*Correspondence:

ioannidis@ece.neu.edu

²Department of Electrical and
Computer Engineering,

Northeastern University, 360

Huntington Avenue, 02115

Boston, MA, USA

Full list of author information is
available at the end of the article

Abstract

Important data mining problems such as nearest-neighbor search and clustering admit theoretical guarantees when restricted to objects embedded in a metric space. Graphs are ubiquitous, and clustering and classification over graphs arise in diverse areas, including, e.g., image processing and social networks. Unfortunately, popular distance scores used in these applications, that scale over large graphs, are not metrics and thus come with no guarantees. Classic graph distances such as, e.g., the chemical distance and the Chartrand-Kubiki-Shultz distance are arguably natural and intuitive, and are indeed also metrics, but they are intractable: as such, their computation does not scale to large graphs. We define a broad family of graph distances, that includes both the chemical and the Chartrand-Kubiki-Shultz distances, and prove that these are all metrics. Crucially, we show that our family includes metrics that are tractable. Moreover, we extend these distances by incorporating auxiliary node attributes, which is important in practice, while maintaining both the metric property and tractability.

Keywords: metric spaces; graph distances; graph matching; graph isomorphism; convex optimization; spectral algorithms

1 Introduction

Graph similarity and the related problem of graph isomorphism have a long history in data mining, machine learning, and pattern recognition [1–3]. *Graph distances* naturally arise in this literature: intuitively, given two (unlabeled) graphs, their distance is a score quantifying their structural differences. A highly desirable property for such a score is that it is a *metric*, i.e., it is non-negative, symmetric, positive-definite, and, crucially, satisfies the triangle inequality. Metrics exhibit significant computational advantages over non-metrics. For example, operations such as nearest-neighbor search [4–6], clustering [7], outlier detection [8], and diameter computation [9] admit fast algorithms precisely when performed over objects embedded in a metric space. To this end, proposing *tractable* graph metrics is of paramount importance in applying such algorithms to graphs.

Unfortunately, graph metrics of interest are often computationally expensive. A well-known example is the *chemical distance* [10]. Formally, given graphs G_A and G_B , represented by their adjacency matrices $A, B \in \{0, 1\}^{n \times n}$, the chemical distance $d_{\mathbb{P}^n}(A, B)$ is defined in terms of a mapping between the two graphs that minimizes their edge discrepancies, i.e.:

$$d_{\mathbb{P}^n}(A, B) = \min_{P \in \mathbb{P}^n} \|AP - PB\|_F, \quad (1)$$

where \mathbb{P}^n is the set of permutation matrices of size n and $\|\cdot\|_F$ is the Frobenius norm (see Sec. 2 for definitions). The *Chartrand-Kubiki-Shultz (CKS)* [11] distance is an alternative: CKS is again given by (1) but, instead of edges, matrices A and B contain the pairwise shortest path distances between any two nodes.

The chemical and CKS distances have important properties. First, they are zero if and only if the graphs are isomorphic, which appeals to both intuition and practice; second, as desired, they are metrics over the quotient space defined by graph isomorphism (see Sec. 2); third, they have a natural interpretation, capturing global structural similarities between graphs. However, finding an optimal permutation P is notoriously hard; graph isomorphism, which is equivalent to deciding if there exists a permutation P such that $AP = PB$ (for both adjacency and path matrices), is famously a problem that is neither known to be in P nor shown to be NP-hard [12]. There is a large and expanding literature on scalable heuristics to estimate the optimal permutation P [13–16]. Despite their computational advantages, unfortunately, using them to approximate $d_{\mathbb{P}^n}(A, B)$ breaks the metric property.

This significantly degrades the performance of many important tasks that rely on computing distances between graphs. For example, there is a clear separation on the approximability of clustering over metric and non-metric spaces [7]. We also demonstrate this empirically in Section 6 (c.f. Fig. 1): attempting to cluster graphs sampled from well-known families based on non-metric distances significantly increases the misclassification rate, compared to clustering using metrics.

An additional issue that arises in practice is that nodes often have attributes not associated with adjacency. For example, in social networks, nodes may contain profiles with a user’s age or gender; similarly, nodes in molecules may be labeled by atomic numbers. Such attributes are not captured by the chemical or CKS distances. However, in such cases, only *label-preserving* permutations P may make sense (e.g., mapping females to females, oxygens to oxygens, etc.). Incorporating attributes while preserving the metric property is thus important from a practical perspective.

1.1 Contributions

We seek generalizations of the chemical and CKS distances that (a) *satisfy the metric property* and (b) are *tractable*: by this, we mean that they can be computed either by solving a convex optimization problem, or by a polynomial time algorithm. Specifically, we study generalizations of (1) of the form:

$$d_S(A, B) = \min_{P \in S} \|AP - PB\| \quad (2)$$

where $S \subset \mathbb{R}^{n \times n}$ is closed and bounded, $\|\cdot\|$ is a matrix norm, and $A, B \in \mathbb{R}^{n \times n}$ are arbitrary real matrices (representing adjacency, path distances, weights, etc.). We make the following contributions:

- We prove sufficient conditions on S and norm $\|\cdot\|$ under which (2) is a *pseudometric*, i.e., a metric over a quotient space defined by equivalence relation $d_S(A, B) = 0$. In particular, we show that d_S is a pseudometric when:
 - (i) $S = \mathbb{P}^n$ and $\|\cdot\|$ is any entry-wise or operator norm;

- (ii) $S = \mathbb{W}^n$, the set of *doubly stochastic* matrices, $\|\cdot\|$ is an arbitrary entry-wise norm, and A, B are symmetric; a modification on d_S extends this result to both operator norms as well as arbitrary matrices (capturing, e.g., directed graphs); and
- (iii) $S = \mathbb{O}^n$, the set of orthogonal matrices, and $\|\cdot\|$ is the operator or entry-wise 2-norm.

We also characterize the corresponding equivalence classes (see Sec. 3). Relaxations (ii) and (iii) are very important from a practical standpoint. For all matrix norms, computing (2) with $S = \mathbb{W}^n$ is tractable, as it is a convex optimization. For $S = \mathbb{O}^n$, (2) is non-convex but is still tractable, as it reduces to a spectral decomposition. This was known for the Frobenius norm [17]; we prove this is also the case for the operator 2-norm.

- We include node attributes in a natural way in the definition of d_S as both *soft* (i.e., penalties in the objective) or *hard* constraints in Eq. (2). Crucially, we do this *without affecting the pseudometric property and tractability*. This allows us to explore label or feature preserving permutations, that incorporate both (a) exogenous node attributes, such as, e.g., user age or gender in a social network, as well as (b) endogenous, structural features of each node, such as its degree or the number of triangles that pass through it. We numerically show that adding these constraints can speed up the computation of d_S .

From an experimental standpoint, we extensively compare our tractable metrics to several existing heuristic approximations. We also demonstrate the tractability of our metrics by parallelizing their execution using the Alternating Direction Method of Multipliers (ADMM) [18], which we implement over a compute cluster using Apache Spark [19].

1.2 Related Work

Graph distance (or similarity) scores find applications in varied fields such as in image processing [1], chemistry [10, 20], and social network analysis [2, 3]. Graph distances are easy to define when, contrary to our setting, the correspondence between graph nodes is known, i.e., graphs are *labeled* [3, 21, 22]. Beyond the chemical distance, classic examples of distances between unlabeled graphs are the *edit distance* [23, 24] and the *maximum common subgraph distance* [25, 26], both of which also have versions for labeled graphs. Both are pseudometrics and are hard to compute, while existing heuristics [27, 28] do not satisfy the triangle inequality. The *reaction distance* [29] is also a pseudometric, and is directly related to the chemical distance [10] when edits are restricted to edge additions and deletions. Jain [30] also considers an extension of the chemical distance, limited to the Frobenius norm, that incorporates edge attributes. However, it is not immediately clear how to relax the above pseudometrics [29, 30] to attain tractability, while keeping the pseudometric property.

A pseudometric can also be induced by embedding graphs in a metric space and measuring the distance between embeddings [31–33]. Several works follow such an approach, mapping graphs, e.g., to spaces determined by their spectral decomposition [34–36]. In general, in contrast to our pseudometrics, such approaches are not as discriminative, as embeddings summarize graph structure. Continuous

relaxations of graph isomorphism, both convex and non-convex [15, 17, 37], have found applications in a variety of contexts, including social networks [38], computer vision [39], shape detection [40, 41], and neuroscience [42]. Lyzinski et al. [15] in particular show (both theoretically and experimentally) that a non-convex relaxation is advantageous over one of the relaxations we consider here (namely, d_S with $S = \mathbb{W}^n$, $\|\cdot\| = \|\cdot\|_F$) in recovering the optimal permutation P . They also incorporate features via a trace penalty as we do in Sec. 3.2 (c.f. Eq. (17)). None of the above works however focus on the metric properties of the resulting relaxations, which several fail to satisfy [15, 38, 40–42].

Metrics naturally arise in data mining tasks, including clustering [43, 44], Nearest Neighbour (NN) search [4–6], and outlier detection [8]. Some of these tasks become tractable, or admit formal guarantees, precisely when performed over a metric space. For example, finding the nearest neighbor [4–6] or the diameter of a data-set [9] become polylogarithmic under metric assumptions; similarly, approximation algorithms for clustering (which is NP-hard) rely on metric assumptions, whose absence leads to a deterioration of known bounds [7]. Our search for metrics is motivated by these considerations.

The present paper is an extended version of a paper by the same authors that appeared in the 2018 SIAM International Conference on Data Mining [45], which did not contain any proofs. In addition to the material included in the conference version, the present paper contains (a) proofs of all main theorems, establishing sufficient conditions under which a solution to (2) yields a pseudo-metric, (b) a polynomial-time spectral algorithm for computing (2) over the Stiefel manifold, (c) extensions of our metrics to graphs of unequal sizes, and (d) an extended experiment section.

2 Notation and Preliminaries

We begin by introducing some terminology that we use throughout the paper. A summary of our notation can be found in Table 1.

Graphs. We represent an undirected unweighted graph $G(V, E)$ with node set $V = [n] \equiv \{1, \dots, n\}$ and edge set $E \subseteq [n] \times [n]$ by its *adjacency matrix*, i.e. $A = [a_{i,j}]_{i,j \in [n]} \in \{0, 1\}^{n \times n}$ such that $a_{ij} = a_{ji} = 1$ if and only if $(i, j) \in E$. In particular, A is symmetric, i.e. $A = A^\top$. We denote the set of all real, symmetric matrices by \mathbb{S}^n . *Directed* unweighted graphs are represented by (possibly non-symmetric) binary matrices $A \in \{0, 1\}^{n \times n}$, and *weighted* graphs by real matrices $A \in \mathbb{R}^{n \times n}$.

Matrix Norms. Given a matrix $A = [a_{i,j}]_{i,j \in [n]} \in \mathbb{R}^{n \times n}$ and a $p \in \mathbb{N}_+ \cup \{\infty\}$, its *induced* or *operator p -norm* is defined in terms of the vector p -norm through $\|A\|_p = \sup_{x \in \mathbb{R}^n: \|x\|_p=1} \|Ax\|_p$, while its *entry-wise p -norm* is given by $\|A\|_p = (\sum_{i=1}^n \sum_{j=1}^n |a_{ij}|^p)^{1/p}$, for $p \in \mathbb{N}_+$, and $\|A\|_\infty = \max_{i,j} |a_{i,j}|$. We denote the entry-wise 2-norm (i.e., the *Frobenius* norm) as $\|\cdot\|_F$. Recall that all matrix norms on $\mathbb{R}^{n \times n}$ are *equivalent*: given two norms $\|\cdot\|, \|\cdot\|'$, there exist constants $c_1, c_2 > 0$ such that

$$c_1 \|A\|' \leq \|A\| \leq c_2 \|A\|' \quad \text{for all } A \in \mathbb{R}^{n \times n}. \quad (3)$$

Permutation, Doubly Stochastic, and Orthogonal Matrices. We denote the set of *permutation matrices* as

$$\mathbb{P}^n \equiv \{P \in [0, 1]^{n \times n} : P\mathbf{1} = \mathbf{1}, P^\top \mathbf{1} = \mathbf{1}\}, \quad (4)$$

the set of *doubly-stochastic matrices* (i.e., the *Birkhoff polytope*) as

$$\mathbb{W}^n \equiv \{W \in [0, 1]^{n \times n} : W\mathbf{1} = \mathbf{1}, W^\top \mathbf{1} = \mathbf{1}\}, \quad (5)$$

and the set of *orthogonal matrices* (i.e., the *Stiefel manifold*) as

$$\mathbb{O}^n \equiv \{U \in \mathbb{R}^{n \times n} : UU^\top = U^\top U = I\}. \quad (6)$$

Note that

$$\mathbb{P}^n = \mathbb{W}^n \cap \mathbb{O}^n, \quad (7)$$

i.e., permutation matrices are precisely the matrices that are both doubly stochastic and orthogonal. Moreover, the Birkhoff-von Neumann Theorem [46] states that

$$\begin{aligned} \mathbb{W}^n &= \text{conv}(\mathbb{P}^n) \\ &= \left\{ W \in [0, 1]^{n \times n} : W = \sum_{P \in \mathbb{P}^n} \theta_P P, \text{ for some } \theta \in \mathbb{R}_+^{|\mathbb{P}^n|} \text{ s.t. } \theta^\top \mathbf{1} = 1 \right\}, \end{aligned} \quad (8)$$

i.e., the Birkhoff polytope is the convex hull of \mathbb{P}^n . Hence, every doubly stochastic matrix can be written as a convex combination of permutation matrices.

Metrics. Given a set Ω , a function $d : \Omega \times \Omega \rightarrow \mathbb{R}$ is called a *metric*, and the pair (Ω, d) is called a *metric space*, if for all $x, y, z \in \Omega$:

$$d(x, y) \geq 0 \quad (\text{non-negativity}) \quad (9a)$$

$$d(x, y) = 0 \text{ iff } x = y \quad (\text{pos. definiteness}) \quad (9b)$$

$$d(x, y) = d(y, x) \quad (\text{symmetry}) \quad (9c)$$

$$d(x, y) \leq d(x, z) + d(z, y) \quad (\text{triangle inequality}) \quad (9d)$$

A function d is called a *pseudometric* if it satisfies (9a), (9c), and (9d), but the positive definiteness property (9b), also known as the *identity of indiscernibles*, is replaced by the (weaker) property:

$$d(x, x) = 0 \text{ for all } x \in \Omega. \quad (9e)$$

If d is a pseudometric, then $d(x, y) = 0$ defines an equivalence relation $x \sim_d y$ over Ω . A pseudometric is then a metric over Ω / \sim_d , the quotient space of \sim_d .

A function d that satisfies (9a), (9b), and (9d) *but not* the symmetry property (9c) is called a *quasimetric*. If d is a quasimetric, then its *symmetric extension* $\bar{d} : \Omega \times \Omega \rightarrow \mathbb{R}$, defined as

$$\bar{d}(x, y) = d(x, y) + d(y, x), \quad (10)$$

is a metric over Ω .

Graph Isomorphism, Chemical Distance, and CKS Distance. Let $A, B \in \mathbb{R}^{n \times n}$ be the adjacency matrices of two graphs G_A and G_B . Then, G_A and G_B are *isomorphic* if and only if there exists $P \in \mathbb{P}^n$ s.t.

$$P^\top AP = B \quad \text{or, equivalently,} \quad AP = PB. \quad (11)$$

The *chemical distance*, given by (1), extends the second relationship in (11) to capture distances between graphs. More generally, let $\|\cdot\|$ be a matrix norm in $\mathbb{R}^{n \times n}$. For some $\Omega \subseteq \mathbb{R}^{n \times n}$, define $d_S : \Omega \times \Omega \rightarrow \mathbb{R}_+$ as:

$$d_S(A, B) = \min_{P \in S} \|AP - PB\|, \quad (12)$$

where $S \subset \mathbb{R}^{n \times n}$ is a closed and bounded set, so that the infimum is indeed attained.

Note that d_S is the chemical distance (1) when $\Omega = \mathbb{R}^{n \times n}$, $S = \mathbb{P}^n$ and $\|\cdot\| = \|\cdot\|_F$. The Chartrand-Kibiki-Shultz (CKS) distance [11] can also be defined in terms of (12), with matrices A, B containing pairwise path distances between any two nodes; equivalently, CKS is the chemical distance of two weighted complete graphs with path distances as edge weights.

The Weisfeiler-Lehman (WL) Algorithm. The WL algorithm [47] is a heuristic for solving the graph isomorphism problem. We use this algorithm to (a) describe the quotient space over which (12) is a metric when $S = \mathbb{W}^n$ (see Sec. 3), and (b) to generate node embeddings in our experiments (see Sec. 6).

To gain some intuition on the algorithm, note that two isomorphic graphs must have the same degree distribution. More broadly, the distributions of k -hop neighborhoods in the two graphs must also be identical. Building on this observation, to test if two undirected, unweighted graphs are isomorphic, WL *colors* the nodes of a graph $G(V, E)$ iteratively. At iteration 0, each node $v \in V$ receives the same *color* $c^0(v) := 1$. Colors at iteration $k + 1 \in \mathbb{N}$ are defined recursively via

$$c^{k+1}(v) := \text{hash}\left(\text{sort}\left(\text{clist}_v^k\right)\right) \quad (13)$$

where hash is a perfect hash function, and

$$\text{clist}_v^k = [c^k(u) : (u, v) \in E] \quad (14)$$

is a list containing the colors of all of v 's neighbors at iteration k .

Intuitively, two nodes in V share the same color after k iterations if their k -hop neighborhoods are isomorphic. WL terminates when the partition of V induced by colors is stable from one iteration to the next. This coloring extends to weighted directed graphs by appending weights and directions to colors in clist_v^k .

After coloring two graphs G_A, G_B , WL declares a non-isomorphism if their color distributions differ. If not, then they *may* be isomorphic and WL gives a set of *constraints* on candidate isomorphisms: a permutation P under which $AP = PB$ must map nodes in G_A to nodes in G_B of the same color.

3 Main Results

Motivated by the chemical and CKS distances, we establish general conditions on S and $\|\cdot\|$ under which d_S is a metric over Ω , for arbitrary weighted graphs. For concreteness, we focus here on distances between graphs of equal size. Extensions to graphs of unequal size are described in Sec. 5.

3.1 A Family of Graph Metrics

Optimization over Permutation Matrices. Our first result establishes that $d_{\mathbb{P}^n}$ is a pseudometric over *all* weighted graphs when $\|\cdot\|$ is an *arbitrary* entry-wise or operator norm.

Theorem 1 *If $S = \mathbb{P}^n$ and $\|\cdot\|$ is an arbitrary entry-wise or operator norm, then d_S given by (12) is a pseudometric over $\Omega = \mathbb{R}^{n \times n}$.*

Hence, $d_{\mathbb{P}^n}$ is a pseudometric under any entry-wise or operator norm over arbitrary directed, weighted graphs.

Optimization over the Birkhoff Polytope. Our second result states that the pseudometric property extends to the *relaxed* version of the chemical distance, in which permutations are replaced by doubly stochastic matrices.

Theorem 2 *If $S = \mathbb{W}^n$ and $\|\cdot\|$ is an arbitrary entry-wise norm, then d_S given by (12) is a pseudometric over $\Omega = \mathbb{S}^{n \times n}$. If $\|\cdot\|$ is an arbitrary entry-wise or operator norm, then its symmetric extension $\bar{d}_S(A, B) = d_S(A, B) + d_S(B, A)$ is a pseudometric over $\Omega = \mathbb{R}^{n \times n}$.*

Hence, if $S = \mathbb{W}^n$ and $\|\cdot\|$ is an arbitrary entry-wise norm, then (12) defines a pseudometric over *undirected* graphs. The symmetry property (9c) breaks if $\|\cdot\|$ is an operator norm or graphs are directed. In both of these two cases d_S is a quasi-metric over the quotient space Ω/\sim_d , and symmetry is attained via the symmetric extension \bar{d}_S .

Theorem 2 has significant practical implications. In contrast to $d_{\mathbb{P}^n}$ and its extensions implied by Theorem 1, computing $d_{\mathbb{W}^n}$ under any operator or entry-wise norm is tractable, in the sense that involves minimizing a convex function subject to linear constraints [48–50].

Optimization over the Stiefel Manifold. A more limited result applies to the case when S is the Stiefel manifold \mathbb{O}^n :

Theorem 3 *If $S = \mathbb{O}^n$ and $\|\cdot\|$ is either the operator (i.e., spectral) or the entry-wise (i.e., Frobenius) 2-norm, then d_S given by (12) is a pseudometric over $\Omega = \mathbb{R}^{n \times n}$.*

Though (12) is not a convex problem when $S = \mathbb{O}^n$, it is also tractable. Umeyama [17] shows that the optimization can be solved exactly when $\|\cdot\| = \|\cdot\|_F$ and $\Omega = \mathbb{S}^n$ (i.e., for undirected graphs) by performing a spectral decomposition on A and B . We extend this result, showing that the same procedure also applies when $\|\cdot\|$ is the operator 2-norm (see Thm. 7 in Sec. 4.3). In the general case of directed graphs, (12) is a classic example of a problem that can be solved through optimization on manifolds [51].

Equivalence Classes. Observe that the equivalence of matrix norms, as stated by Eq. (3), implies that if $d_S(A, B) = 0$ for one matrix norm $\|\cdot\|$ in (12), it will be so for all. As a result, pseudometrics d_S defined through (12) for a given S have the same quotient space Ω/\sim_{d_S} , irrespectively of norm $\|\cdot\|$. We therefore turn our attention to characterizing this quotient space in the three cases when S is the set of permutation, doubly stochastic, and orthonormal matrices.

When $S = \mathbb{P}^n$, $\Omega/\sim_{d_{\mathbb{P}^n}}$ is the quotient space defined by graph isomorphism: any two adjacency matrices $A, B \in \mathbb{R}^{n \times n}$ satisfy $d_{\mathbb{P}^n}(A, B) = 0$ if and only if their (possibly weighted) graphs are isomorphic.

When $S = \mathbb{W}^n$, the quotient space $\Omega/\sim_{d_{\mathbb{W}^n}}$ has a connection to the Weisfeiler-Lehman (WL) algorithm [47] described in Section 2: Ramana et al. [52] show that $d_{\mathbb{W}^n}(A, B) = 0$ if and only if G_A and G_B receive identical colors by the WL algorithm (see also [53] for another characterization of this quotient space). This equivalence relation is sometimes called fractional linear isomorphism [52].

Finally, if $S = \mathbb{O}^n$ and $\Omega = \mathbb{S}^n$, i.e., graphs are undirected, then $\Omega/\sim_{d_{\mathbb{O}^n}}$ is determined by *co-spectrality*: $d_{\mathbb{O}^n}(A, B) = 0$ if and only if A, B have the same spectrum. When $\Omega = \mathbb{R}^{n \times n}$, $d_{\mathbb{O}^n}(A, B) = 0$ implies that A, B are co-spectral, but co-spectral matrices A, B do not necessarily satisfy $d_{\mathbb{O}^n}(A, B) = 0$. Put differently, the quotient space $\Omega/\sim_{d_{\mathbb{O}^n}}$ in this case is a refinement of the quotient space of co-spectrality.

3.2 Incorporating Metric Embeddings

We have seen that the chemical distance $d_{\mathbb{P}^n}$ can be *relaxed* to $d_{\mathbb{W}^n}$ or $d_{\mathbb{O}^n}$, gaining tractability while still maintaining the metric property. In practice, nodes in a graph often contain additional attributes that one might wish to leverage when computing distances. In this section, we show that such attributes can be seamlessly incorporated in d_S either as soft or hard constraints, *without violating the metric property*.

Metric Embeddings. Given a graph G_A of size n , a *metric embedding* of G_A is a mapping $\psi_A : [n] \rightarrow \tilde{\Omega}$ from the nodes of the graph to a metric space $(\tilde{\Omega}, \tilde{d})$. That is, ψ_A maps nodes of the graph to $\tilde{\Omega}$, where $\tilde{\Omega}$ is endowed with a metric \tilde{d} . We refer to a graph endowed with an embedding ψ_A as an *embedded graph*, and denote this by (A, ψ_A) , where $A \in \mathbb{R}^{n \times n}$ is the adjacency matrix of G_A . We list two examples:

Example 1: Node Attributes. Consider an embedding of a graph to $(\mathbb{R}^k, \|\cdot\|_2)$ in which every node $v \in V$ is mapped to a k -dimensional vector describing “local” attributes. These can be *exogenous*: e.g., features extracted from a user’s profile (age, binarized gender, etc.) in a social network. Alternatively, attributes may be *endogenous* or *structural*, extracted from the adjacency matrix A , e.g., the node’s degree, the size of its k -hop neighborhood, its page-rank, etc.

Example 2: Node Colors. Let $\tilde{\Omega}$ be an arbitrary finite set endowed with the Kronecker delta as a metric, that is, for $s, s' \in \tilde{\Omega}$,

$$\tilde{d}(s, s') = \begin{cases} 0, & \text{if } s = s' \\ \infty, & \text{o.w.} \end{cases} \quad (15)$$

Given a graph G_A , a mapping $\psi_A : [n] \rightarrow \tilde{\Omega}$ is then a metric embedding. The values of $\tilde{\Omega}$ are invariably called *colors* or *labels*, and a graph embedded in $\tilde{\Omega}$ is a *colored* or *labeled* graph. Colors can again be *exogenous* or *structural*: e.g., if the graph represents an organic molecule, colors can correspond to atoms, while structural colors can be, e.g., the output of the WL algorithm (see Section 2) after k iterations.

As discussed below, node attributes translate to *soft* constraints in metric (12), while node colors correspond to *hard* constraints. The unified view through embeddings allows us to establish metric properties for both simultaneously (c.f. Theorems 4 and 5).

Embedding Distance. Consider two embedded graphs $(A, \psi_A), (B, \psi_B)$ of size n that are embedded in the same metric space $(\tilde{\Omega}, \tilde{d})$. For $u \in [n]$, a node in the first graph, and $v \in [n]$, a node in the second graph, the embedded distance between the two nodes is given by $\tilde{d}(\psi_A(u), \psi_B(v))$. Let $D_{\psi_A, \psi_B} = [\tilde{d}(\psi_A(u), \psi_B(v))]_{u \in V, v \in V} \in \mathbb{R}_+^{n \times n}$ be the corresponding matrix of embedded distances. After mapping nodes to the same metric space, it is natural to seek $P \in \mathbb{P}^n$ that preserve the *embedding distance*. This amounts to finding a $P \in \mathbb{P}^n$ that minimizes:

$$\text{tr}(P^\top D_{\psi_A, \psi_B}) = \sum_{u, v \in [n]} P_{u, v} \tilde{d}(\psi_A(u), \psi_B(v)). \quad (16)$$

Note that, in the case of colored graphs and the Kronecker delta distance, minimizing (16) finds a $P \in \mathbb{P}^n$ that maps nodes in A to nodes in B of equal color. It is not hard to verify that $\min_{P \in \mathbb{P}^n} \text{tr}(P^\top D_{\psi_A, \psi_B})$ induces a metric between graphs embedded in $(\tilde{\Omega}, \tilde{d})$. In fact, this follows from the more general theorem we prove below (Theorem. 4) for $A = B = 0$, i.e., for distances between embedded graphs with no edges.

Despite the combinatorial nature of \mathbb{P}^n , the problem of minimizing (16) over \mathbb{P}^n is a maximum weighted matching problem, which can be solved through, e.g., the Hungarian algorithm [54], in $O(n^3)$ time [55]. We note that this metric is not as expressive as (12): depending on the definition of the embeddings ψ_A, ψ_B , attributes may only capture “local” similarities between nodes, as opposed to the “global” view of a mapping attained by (12).

A Unified, Tractable Metric. Motivated by the above considerations, we focus on unifying the “global” metric (12) with the “local” metrics induced by arbitrary graph embeddings. Given a metric space $(\tilde{\Omega}, \tilde{d})$, let $\Psi_{\tilde{\Omega}}^n = \{\psi : [n] \rightarrow \tilde{\Omega}\}$ be the set of all mappings from $[n]$ to $\tilde{\Omega}$. Then, given two embedded graphs $(A, \psi_A), (B, \psi_B) \in \mathbb{R}^{n \times n} \times \Psi_{\tilde{\Omega}}^n$, we define:

$$d_S((A, \psi_A), (B, \psi_B)) = \min_{P \in S} [\|AP - PB\| + \text{tr}(P^\top D_{\psi_A, \psi_B})] \quad (17)$$

for some compact set $S \subset \mathbb{R}^{n \times n}$ and matrix norm $\|\cdot\|$. Our next result states that incorporating this linear term does not affect the pseudometric property of d_S .

Theorem 4 *If $S = \mathbb{P}^n$ and $\|\cdot\|$ is an arbitrary entry-wise or operator norm, then d_S given by (17) is a pseudometric over the set of embedded graphs $\Omega = \mathbb{R}^{n \times n} \times \Psi_{\tilde{\Omega}}^n$.*

We stress here that this result is non-obvious: it is not true that adding *any* linear term to d_S leads to a quantity that satisfies the triangle inequality. It is precisely because D_{ψ_A, ψ_B} contains pairwise distances that Theorem 4 holds. We can similarly extend Theorem 2:

Theorem 5 *If $S = \mathbb{W}^n$ and $\|\cdot\|$ is an arbitrary entry-wise norm, then d_S given by (17) is a pseudometric over $\Omega = \mathbb{S}^n \times \Psi_\Omega^n$, the set of symmetric graphs embedded in $(\tilde{\Omega}, \tilde{d})$. Moreover, if $\|\cdot\|$ is an arbitrary entry-wise or operator norm, then the symmetric extension \tilde{d}_S of (17) is a pseudometric over $\Omega = \mathbb{R}^{n \times n} \times \Psi_\Omega^n$.*

Adding the linear term (16) in d_S has significant practical advantages. Beyond expressing exogenous attributes, a linear term involving colors, combined with a Kronecker distance, translates into *hard* constraints: any permutation attaining a finite objective value *must* map nodes in one graph to nodes of the same color. Theorem 5 therefore implies that such constraints can thus be added to the optimization problem, while maintaining the metric property. In practice, as the number of variables in optimization problem (12) is n^2 , incorporating such hard constraints can significantly reduce the problem's computation time; we illustrate this in Section 6. Note that adding (16) to $d_{\mathbb{O}^n}$ does *not* preserve the metric property.

4 Proofs of Main Results

4.1 Proof of Theorems 1–3.

We define several properties that play a crucial role in our proofs.

Definition 1 *We say that a set $S \subseteq \mathbb{R}^{n \times n}$ is closed under multiplication if $P, P' \in S$ implies that $P \cdot P' \in S$.*

Definition 2 *We say that $S \subseteq \mathbb{R}^{n \times n}$ is closed under transposition if $P \in S$ implies that $P^T \in S$, and closed under inversion if $P \in S$ implies that $P^{-1} \in S$.*

Definition 3 *Given a matrix norm $\|\cdot\|$, we say that set $S \subseteq \mathbb{R}^{n \times n}$ is contractive w.r.t. $\|\cdot\|$ if $\|AP\| \leq \|A\|$ and $\|PA\| \leq \|A\|$, for all $P \in S$ and $A \in \mathbb{R}^{n \times n}$. Put differently, S is contractive if and only if every linear transform $P \in S$ is a contraction w.r.t. $\|\cdot\|$.*

The proofs of Theorems 1–3 rely on several common lemmas. The first three establish conditions under which (12) satisfies the triangle inequality (9d), symmetry (9c), and weak property (9e), respectively:

Lemma 1 *Given a matrix norm $\|\cdot\|$, suppose that set $S \subseteq \mathbb{R}^{n \times n}$ is (a) contractive w.r.t. $\|\cdot\|$, and (b) closed under multiplication. Then, for any $A, B, C \in \mathbb{R}^{n \times n}$, d_S given by (12) satisfies $d_S(A, C) \leq d_S(A, B) + d_S(B, C)$.*

Proof Consider $P' \in \arg \min_{P \in S} \|AP - PB\|$, and $P'' \in \arg \min_{P \in S} \|BP - PC\|$. Then, from closure under multiplication, $P'P'' \in S$. Hence,

$$\begin{aligned} d_S(A, C) &\leq \|AP'P'' - P'P''C\| = \|AP'P'' - P'BP'' + P'BP'' - P'P''C\| \\ &\leq \|AP'P'' - P'BP''\| + \|P'BP'' - P'P''C\| \\ &= \|(AP' - P'B)P''\| + \|P'(BP'' - P''C)\| \\ &\leq \|AP' - P'B\| + \|BP'' - P''C\| = d_S(A, B) + d_S(B, C), \end{aligned}$$

where the last inequality follows from the fact that P', P'' are contractions. \square

Lemma 2 *Given a matrix norm $\|\cdot\|$, suppose that $S \subset \mathbb{R}^{n \times n}$ is (a) contractive w.r.t. $\|\cdot\|$, and (b) closed under inversion. Then, for all $A, B \in \mathbb{R}^{n \times n}$, $d_S(A, B) = d_S(B, A)$.*

Proof Observe that property (b) implies that, for all $P \in S$, P is invertible and $P^{-1} \in S$. Hence,

$$\|AP - PB\| = \|PP^{-1}AP - PBP^{-1}P\| = \|P(P^{-1}A - BP^{-1})P\| \leq \|BP^{-1} - P^{-1}A\|,$$

as P is a contraction w.r.t. $\|\cdot\|$. We can similarly show that $\|BP^{-1} - P^{-1}A\| \leq \|AP - PB\|$, hence $\|AP - PB\| = \|BP^{-1} - P^{-1}A\|$. As S is closed under inversion,

$$\min_{P \in S} f(P) = \min_{P: P^{-1} \in S} f(P),$$

for every $f : S \rightarrow \mathbb{R}$. Hence

$$\begin{aligned} d_S(A, B) &= \min_{P \in S} \|BP^{-1} - P^{-1}A\| = \min_{P: P^{-1} \in S} \|BP^{-1} - P^{-1}A\| \\ &= \min_{P \in S} \|BP - PA\| = d_S(B, A). \end{aligned} \quad \square$$

Lemma 3 *If $I \in S$, then $d_S(A, A) = 0$ for all $A \in \mathbb{R}^{n \times n}$.*

Proof Indeed, if $I \in S$, then $0 \leq d_S(A, A) \leq \|AI - IA\| = 0$. \square

Both the set of permutation matrices \mathbb{P}^n and the Stiefel manifold \mathbb{O}^n are groups w.r.t. matrix multiplication: they are closed under multiplication, contain the identity I , and are closed under inversion. Hence, if they are also contractive w.r.t. a matrix norm $\|\cdot\|$, $d_{\mathbb{P}^n}$ and $d_{\mathbb{O}^n}$ defined in terms of this norm satisfy all assumptions of Lemmas 1–3. We therefore turn our attention to this property.

Lemma 4 *Let $\|\cdot\|$ be any operator or entry-wise norm. Then, $S = \mathbb{P}^n$ is contractive w.r.t. $\|\cdot\|$.*

Proof Observe first that all vector p -norms are invariant to permutations of a vector's entries; hence, for any vector $x \in \mathbb{R}^d$, if $P \in \mathbb{P}^n$, $\|Px\|_p = \|x\|_p$. Hence, if $\|\cdot\|$ is an operator p -norm, $\|P\| = 1$, for all $P \in S$. Every operator norm is *submultiplicative*; as a result $\|PA\| \leq \|P\|\|A\| = \|A\|$ and, similarly, $\|AP\| \leq \|A\|$, so the lemma follows

for operator norms. On the other hand, if $\|\cdot\|$ is an entry-wise norm, then $\|A\|$ is invariant to permutations of either A 's rows or columns. Matrices PA and AP precisely amount to such permutations, so $\|PA\| = \|AP\| = \|A\|$ and the lemma follows also for entrywise norms. \square

Hence, Theorem 1 follows as a direct corollary of Lemmas 1–4. Indeed, $d_{\mathbb{P}^n}$ is non-negative, symmetric by Lemmas 2 and 4, satisfies the triangle inequality by Lemmas 1 and 4, as well as property (9e) by Lemma 3; hence $d_{\mathbb{P}^n}$ is a pseudometric over $\mathbb{R}^{n \times n}$. Our next lemma shows that the Stiefel manifold \mathbb{O}^n is contractive for 2-norms:

Lemma 5 *Let $\|\cdot\|$ be the operator (i.e., spectral) or the entry-wise (i.e., Frobenius) 2-norm. Then, $S = \mathbb{O}^n$ is contractive w.r.t. $\|\cdot\|$.*

Proof Any $U \in \mathbb{O}^n$ is an orthogonal matrix; hence, $\|U\|_2 = \|U\|_F = 1$. Both norms are submultiplicative: the first as an operator norm, the second from the Cauchy-Schwartz inequality. Hence, for $U \in \mathbb{O}^n$, we have $\|UA\| \leq \|U\| \|A\| = \|A\|$.

Note that an alternative proof can be obtained by the fact that both norms are unitarily invariant (see Lemma 12). \square

Theorem 3 therefore follows from Lemmas 1–3 and Lemma 5, along with the fact that \mathbb{O}^n is a group. Note that \mathbb{O}^n is *not* contractive w.r.t. other norms, e.g., $\|\cdot\|_1$ or $\|\cdot\|_\infty$.

To prove Theorem 2, we first show that Lemma 4 along with the Birkhoff-von Neumann theorem imply that \mathbb{W}^n is also contractive:

Lemma 6 *Let $\|\cdot\|$ be any operator or entry-wise norm. Then, \mathbb{W}^n is contractive w.r.t. $\|\cdot\|$.*

Proof By the Birkhoff-von Neumann theorem [46],

$$\mathbb{W}^n = \text{conv}(\mathbb{P}^n).$$

Hence, for any $W \in \mathbb{W}^n$ there exist $P_i \in \mathbb{P}^n$, $\theta_i > 0$, $i = 1, \dots, k$, such that $W = \sum_{i=1}^k \theta_i P_i$ and $\sum_{i=1}^k \theta_i = 1$. Both operator and entrywise p -norms are convex functions; hence, for any $A \in \mathbb{R}^{n \times N}$:

$$\begin{aligned} \|WA\| &= \left\| \sum_{i=1}^k \theta_i P_i A \right\| \leq \sum_{i=1}^k \theta_i \|P_i A\|, && \text{by Jensen's inequality,} \\ &\leq \sum_{i=1}^k \theta_i \|A\|, && \text{by Lemma 4,} \\ &= \|A\| \end{aligned}$$

The statement $\|AW\| \leq \|A\|$ follows similarly. \square

Unfortunately, the Birkhoff polytope \mathbb{W}^n is *not* a group, as it is not closed under inversion. Nevertheless, it is closed under transposition; in establishing (partial) symmetry of $d_{\mathbb{W}^n}$, we leverage the following lemma:

Lemma 7 *Suppose that $\|\cdot\|$ is transpose-invariant, and $S \subseteq \mathbb{R}^{n \times n}$ is closed under transposition. Then, $d_S(A, B) = d_S(B, A)$ for all $A, B \in \mathbb{S}^n$.*

Proof By transpose invariance and the symmetry of A and B , we have that:

$$\|AP - PB\| = \|BP^\top - P^\top A\|.$$

Moreover, as S is closed under transposition, for every $f : S \rightarrow \mathbb{R}$,

$$\min_{P \in S} f(P) = \min_{P: P^\top \in S} f(P).$$

Hence, $d_S(A, B) = \min_{P \in S} \|BP^\top - P^\top A\| = \min_{P: P^\top \in S} \|BP^\top - P^\top A\| = d_S(B, A)$. \square

The first part of Theorem 2 therefore follows from Lemmas 1, 3, 6, and 7: this is because \mathbb{W}^n contains the identity I and is closed under both transposition and multiplication, while all entry-wise norms are transpose invariant.

To prove the second part, observe that operator norms are not transpose invariant. However, if $\|\cdot\|$ is an operator norm, or $\Omega = \mathbb{R}^{n \times n}$, then Lemma 6 and Lemma 1 imply that $d_{\mathbb{W}^n}$ satisfies non-negativity (9a) and the triangle inequality (9d), while Lemma 3 implies that it satisfies (9e). These properties are inherited by extension \bar{d}_S , given by (10), which also satisfies symmetry (9c), and the second part of Theorem 2 follows.

4.2 Proof of Theorems 4 and 5

We begin by establishing conditions under which d_S satisfies the triangle inequality (9d). We note that, in contrast to Lemma 1, we require the additional condition that $S \subseteq \mathbb{W}^n$, which is not satisfied by \mathbb{O}^n .

Lemma 8 *Given a norm $\|\cdot\|$, suppose that $S \subseteq \mathbb{R}^{n \times n}$ is (a) contractive w.r.t. $\|\cdot\|$, (b) closed under multiplication, and (c) is a subset of \mathbb{W}^n , i.e., contains only doubly stochastic matrices. Then, for any $(A, \psi_A), (B, \psi_B), (C, \psi_C)$ in $\mathbb{R}^{n \times n} \times \Psi_\Omega$,*

$$d_S((A, \psi_A), (C, \psi_B)) \leq d_S((A, \psi_A), (B, \psi_B)) + d_S((B, \psi_B), (C, \psi_C)).$$

Proof Consider

$$P' \in \arg \min_{P \in S} \left(\|AP - PB\| + \text{tr} \left(P^\top D_{\psi_A, \psi_B} \right) \right),$$

and

$$P'' \in \arg \min_{P \in S} \left(\|BP - PC\| + \text{tr} \left(P^\top D_{\psi_B, \psi_C} \right) \right).$$

Then, from closure under multiplication, $P'P'' \in S$. We have that

$$d_S((A, \psi_A), (C, \psi_C)) \leq \|AP'P'' - P'P''C\| + \text{tr} \left[(P'P'')^\top D_{\psi_A, \psi_C} \right]$$

As in the proof of Lemma 1, we can show that

$$\begin{aligned}
\|AP'P'' - P'P''C\| &= \|AP'P'' - P'BP'' + P'BP'' - P'P''C\| \\
&\leq \|AP'P'' - P'BP''\| + \|P'BP'' - P'P''C\| \\
&= \|(AP' - P'B)P''\| + \|P'(BP'' - P''C)\| \\
&\leq \|(AP' - P'B)\| \|P''\| + \|P'\| \|(BP'' - P''C)\| \\
&\leq \|AP' - P'B\| + \|BP'' - P''C\|
\end{aligned}$$

using the fact that both P' and P'' are contractions. On the other hand,

$$\begin{aligned}
\text{tr}[(P'P'')^\top D_{\psi_A\psi_C}] &= \sum_{u,v \in [n]} \sum_{k \in [n]} (P'_{uk} P''_{kv} \tilde{d}(\psi_A(u), \psi_C(v))) \\
&\leq \sum_{u,v \in [n]} \sum_{k \in [n]} [P'_{uk} P''_{kv} (\tilde{d}(\psi_A(u), \psi_B(k)) + \tilde{d}(\psi_B(k), \psi_C(v)))] \\
&\quad (\text{as } \tilde{d} \text{ is a metric, and } P', P'' \text{ are non-negative}) \\
&= \sum_{u,k \in [n]} P'_{uk} \tilde{d}(\psi_A(u), \psi_B(k)) \sum_{v \in [n]} P''_{kv} + \\
&\quad \sum_{k,v \in [n]} P''_{kv} \tilde{d}(\psi_B(k), \psi_C(v)) \sum_{u \in [n]} P'_{uk} \\
&\leq \text{tr}((P')^\top D_{\psi_A, \psi_B}) + \text{tr}((P'')^\top D_{\psi_B, \psi_C}),
\end{aligned}$$

where the last inequality follows as both P, P^\top are $\|\cdot\|_1$ -norm bounded by 1 for every $P \in S$. \square

The weak property (9e) is again satisfied provided the identity is included in S .

Lemma 9 *If $I \in S$, then $d_S((A, \psi_A), (A, \psi_A)) = 0$ for all $A \in \mathbb{R}^{n \times n}$.*

Proof Indeed, $0 \leq d_S((A, \psi_A), (A, \psi_A)) \leq \|AI - IA\| + \sum_{u \in [n]} \tilde{d}(\psi_A(u), \psi_A(u)) = 0$. \square

To attain symmetry over $\Omega = \mathbb{R}^{n \times n} \times \Psi_\Omega^n$, we again rely on closure under inversion, as in Lemma 2; nonetheless, in contrast to Lemma 2, due to the linear term, we also need to assume the orthogonality of elements of S .

Lemma 10 *Given a norm $\|\cdot\|$, suppose that S (a) is contractive w.r.t. $\|\cdot\|$, (b) is closed under inversion, and (c) is a subset of \mathbb{O}^n , i.e., contains only orthogonal matrices. Then, $d_S((A, \psi_A), (B, \psi_B)) = d_S((B, \psi_B), (A, \psi_A))$ for all $(A, \psi_A), (B, \psi_B) \in \mathbb{R}^{n \times n} \times \Psi_\Omega$.*

Proof As in the proof of Lemma 2, we can show that contractiveness w.r.t. $\|\cdot\|$ along with closure under inversion imply that: $\|AP - PB\| = \|BP^{-1} - P^{-1}A\|$. As S is closed under inversion, $\min_{P \in S} f(P) = \min_{P^{-1} \in S} f(P)$ for all $f : S \rightarrow \mathbb{R}$, while

orthogonality implies $P^{-1} = P^\top$ for all $P \in S$. Hence, $d_S((A, \psi_A), (B, \psi_B))$ equals

$$\begin{aligned} \min_{P \in S} \left[\|AP - PB\| + \operatorname{tr} \left(P^\top D_{\psi_A, \psi_B} \right) \right] &= \min_{P \in S} \left[\|BP^{-1} - P^{-1}A\| + \operatorname{tr} \left(P^{-1} D_{\psi_A, \psi_B} \right) \right] \\ &= \min_{P \in S} \left[\|BP^{-1} - P^{-1}A\| + \operatorname{tr} \left((P^{-1})^\top D_{\psi_A, \psi_B}^\top \right) \right] \\ &= \min_{P \in S} \left[\|BP^{-1} - P^{-1}A\| + \operatorname{tr} \left((P^{-1})^\top D_{\psi_B, \psi_A} \right) \right] \\ &= \min_{P: P^{-1} \in S} \left[\|BP^{-1} - P^{-1}A\| + \operatorname{tr} \left((P^{-1})^\top D_{\psi_B, \psi_A} \right) \right] \\ &= d_S((B, \psi_B), (A, \psi_A)). \quad \square \end{aligned}$$

Theorem 4 therefore follows from the above lemmas, as $S = \mathbb{P}^n$ contains I , it is closed under multiplication and inversion, is a subset of $\mathbb{W}^n \cap \mathbb{O}^n$ by (7), and is contractive w.r.t. all operator and entrywise norms. Theorem 5 also follows by using the following lemma, along with Lemmas 8 and 9.

Lemma 11 *Suppose that $\|\cdot\|$ is transpose invariant, and S is closed under transposition. Then, $d_S((A, \psi_A), (B, \psi_B)) = d_S((B, \psi_B), (A, \psi_A))$ for all $(A, \psi_A), (B, \psi_B) \in \mathbb{S}^n \times \Psi_{\Omega}$.*

Proof By the transpose invariance of $\|\cdot\|$ and the symmetry of A and B , we have that: $\|AP - PB\| = \|BP^\top - P^\top A\|$. Moreover, as S is closed under transposition, $\min_{P \in S} f(P) = \min_{P: P^\top \in S} f(P)$ for any $f: S \rightarrow \mathbb{R}$. Hence, $d_S((A, \psi_A), (B, \psi_B))$ equals

$$\begin{aligned} \min_{P \in S} \left[\|AP - PB\| + \operatorname{tr} \left(P^\top D_{\psi_A, \psi_B} \right) \right] &= \min_{P \in S} \left[\|BP^\top - P^\top A\| + \operatorname{tr} \left(P D_{\psi_A, \psi_B}^\top \right) \right] \\ &= \min_{P: P^\top \in S} \|BP^\top - P^\top A\| + \operatorname{tr} \left((P^\top)^\top D_{\psi_B, \psi_A} \right) \\ &= d_S((B, \psi_B), (A, \psi_A)) \quad \square \end{aligned}$$

4.3 Metric Computation Over the Stiefel Manifold.

In this section, we describe how to compute the metric d_S in polynomial time when $S = \mathbb{O}^n$ and $\|\cdot\|$ is the Frobenius norm or the operator 2-norm. The algorithm for the Frobenius norm, and the proof of its correctness, is due to Umeyama [17]; we reprove it for completeness, along with its extension to the operator norm.

Both cases make use of the following lemma:

Lemma 12 *For any matrix $M \in \mathbb{R}^{n \times n}$ and any matrix $P \in \mathbb{O}^n$ we have that $\|PM\| = \|MP\| = \|M\|$, where $\|\cdot\|$ is either the Frobenius or operator 2-norm.*

Proof Recall that the operator 2-norm $\|\cdot\|_2$ is $\|M\|_2 = \sup_{x \neq 0} \|Mx\|_2 / \|x\|_2 = \sqrt{\sigma_{\max}(M^\top M)} = \sqrt{\sigma_{\max}(MM^\top)} = \|M^\top\|_2$. where σ_{\max} denotes the largest singular value. Hence, $\|PM\|_2 = \sup_{x \neq 0} \|PMx\|_2 / \|x\|_2 = \sqrt{\sigma_{\max}(M^\top P^\top PM)} = \sqrt{\sigma_{\max}(M^\top M)} = \|M\|_2$. as $P^\top P = I$. Using the fact that $\|M\|_2 = \|M^\top\|_2$ for all $M \in \mathbb{R}^{n \times n}$, as well as that $PP^\top = I$, we can show that $\|MP\|_2 = \|P^\top M^\top\|_2 = \|M^\top\|_2 = \|M\|_2$.

The Frobenius norm is $\|M\|_F = \sqrt{\operatorname{tr}(M^\top M)} = \sqrt{\operatorname{tr}(MM^\top)} = \|M^\top\|_F$, hence $\|PM\|_F = \sqrt{\operatorname{tr}(M^\top P^\top PM)} = \sqrt{\operatorname{tr}(M^\top M)} = \|M\|_F$ and, as in the case of the operator norm, we can similarly show $\|MP\|_F = \|P^\top M^\top\|_F = \|M^\top\|_F = \|M\|_F$. \square

In both norm cases, for $A, B \in \mathbb{S}^n$, we can compute d_S using a simple spectral decomposition, which dominates computations and can be performed in $O(n^3)$ time. Let $A = U\Sigma_A U^T$ and $B = V\Sigma_B V^T$ be the spectral decomposition of A and B . As A and B are real and symmetric, we can assume $U, V \in \mathbb{O}^n$. Recall that $U^{-1} = U^T$ and $V^{-1} = V^T$, while Σ_A and Σ_B are diagonal and contain the eigenvalues of A and B sorted in increasing order; this ordering matters for computations below.

The following theorem establishes that this decomposition readily yields the distance d_S , as well as the optimal orthogonal matrix P^* , when $\|\cdot\| = \|\cdot\|_F$:

Theorem 6 ([17]) $d_S(A, B) \triangleq \min_{P \in \mathbb{S}} \|AP - PB\|_F = \|\Sigma_A - \Sigma_B\|_F$ and the minimum is attained by $P^* = UV^T$.

Proof The proof makes use of the following lemma by Hoffman and Wielandt [56]:

Lemma 13 ([56]) *If A and B are Hermitian matrices with eigenvalues $a_1 \leq a_2 \leq \dots \leq a_n$ and $b_1 \leq b_2 \leq \dots \leq b_n$ then $\|A - B\|_F^2 \geq \sum_{i=1}^n (a_i - b_i)^2$*

Note that if Σ_A and Σ_B are diagonal matrices with the ordered eigenvalues of A and B in the diagonal, then the conclusion of Lemma 13 can be written as $\|A - B\|_F \geq \|\Sigma_A - \Sigma_B\|_F$. For any $P \in \mathbb{O}^n$ and $\|\cdot\| = \|\cdot\|_F$ we have

$$\begin{aligned} \|AP - PB\| &= \|(A - PBP^{-1})P\| \stackrel{\text{Lem. 12}}{=} \|A - PBP^T\| = \|U\Sigma_A U^T - PV\Sigma_B V^T P^T\| \\ &= \|U(\Sigma_A - U^T P V \Sigma_B V^T P^T U)U^T\| \stackrel{\text{Lem. 12}}{=} \|\Sigma_A - U^T P V \Sigma_B V^T P^T U\| \\ &= \|\Sigma_A - \Delta \Sigma_B \Delta^T\| \end{aligned}$$

where we define $\Delta \equiv U^T P V$. As a product of orthogonal matrices, $\Delta \in \mathbb{O}^n$. Notice that

$$\begin{aligned} \|\Sigma_A - \Delta \Sigma_B \Delta^T\| &= \|\Sigma_A - \Delta \Sigma_A \Delta^T + \Delta(\Sigma_B - \Sigma_A)\Delta^T\| \leq \|\Sigma_A - \Delta \Sigma_A \Delta^T\| + \|\Delta(\Sigma_B - \Sigma_A)\Delta^T\| \\ &\stackrel{\text{Lem. 12}}{=} \|\Sigma_A - \Delta \Sigma_A \Delta^T\| + \|\Sigma_B - \Sigma_A\|. \end{aligned}$$

Therefore, for any $P \in \mathbb{O}^n$, $\|\Sigma_A - \Sigma_B\| \leq d_S(A, B) \leq \|\Sigma_A - \Delta \Sigma_A \Delta^T\| + \|\Sigma_B - \Sigma_A\|$, where the first inequality follows by Lemma 13 if we notice that $\|AP - PB\| = \|A - PBP^{-1}\|$ and that PBP^{-1} and B have the same spectrum for any P . If we choose $P = UV^T$ then $\Delta = I$ and the result follows. \square

We can compute d_S when $S = \mathbb{O}^n$ and $\|\cdot\|$ is the operator norm in the exact same way.

Theorem 7 *Let $\|\cdot\| = \|\cdot\|_2$ be the operator 2-norm. Then, $d_S(A, B) \triangleq \min_{P \in \mathbb{S}} \|AP - PB\|_2 = \|\Sigma_A - \Sigma_B\|_2$ and the minimum is attained by $P^* = UV^T$.*

Proof The proof follows the same steps as the proof of Theorem 6, using Lemma 14 below instead of Lemma 13.

Lemma 14 *If A and B are Hermitian matrices with eigenvalues $a_1 \leq a_2 \leq \dots \leq a_n$ and $b_1 \leq b_2 \leq \dots \leq b_n$ then $\|A - B\|_2 \geq \max_i |a_i - b_i|$.*

Proof This is the second exercise following Corollary 6.3.4 in Horn and Johnson [57]. We reprove this here for completeness.

Let $\tilde{B} = -B$ have eigenvalues $\tilde{b}_1 \leq \tilde{b}_2 \leq \dots \leq \tilde{b}_n$ and let $C = A + \tilde{B}$ have eigenvalues $c_1 \leq c_2 \leq \dots \leq c_n$. We make use of the following lemma by Weyl (see Theorem 4.3.1 (Weyl), page 239, in [57]) to lower-bound c_n .

Lemma 15 (Weyl) *If X and Y are Hermitian with eigenvalues $x_1 \leq \dots \leq x_n$ and $y_1 \leq \dots \leq y_n$ and if $X + Y$ has eigenvalues $w_1 \leq \dots \leq w_n$ then $x_{i-j+1} + y_j \leq w_i$ for all $i = 1, \dots, n$ and $j = 1, \dots, i$.*

If we choose $X = \tilde{B}$, $Y = A$ and $i = n$ we get $a_j + \tilde{b}_{n+1-j} \leq c_n$ for all $j = 1, \dots, n$. Since $\tilde{b}_{n+1-j} = -b_j$ we get that $a_j - b_j \leq c_n$, for any j . Similarly, by exchanging the role of A and B , we can lower bound the largest eigenvalue of $B - A$, say d_n , by $b_j - a_j$ for any j . Notice that, by definition of the operator norm and the fact that $A - B$ is Hermitian, $\|A - B\|_2 \geq |c_n|$ and $\|B - A\|_2 \geq |d_n|$. Since $\|B - A\|_2 = \|A - B\|_2$ we have that $\|A - B\|_2 \geq \max\{|c_n|, |d_n|\} \geq \max\{c_n, d_n\} \geq \max\{a_j - b_j, b_j - a_j\} = |a_j - b_j|$ for all j . Taking the maximum over j we get that $\|A - B\|_2 \geq \max_j |a_j - b_j|$, and the lemma follows. \square

Note again that if Σ_A and Σ_B are diagonal matrices with the ordered eigenvalues of A and B in the diagonal, then the conclusion of Lemma 14 can be written as $\|A - B\|_2 \geq \|\Sigma_A - \Sigma_B\|_2$. The proof of Thm. 7 proceeds along the same steps as the above proof, using again the fact that, by Lemma 12, $\|M\|_2 = \|MP\|_2 = \|PM\|_2$ for any $P \in \mathbb{O}^n$ and any matrix M , along with Lemma 15. \square

5 Graphs of Different Sizes

For simplicity, we have described distances over graphs of equal sizes. There are several applications [15, 58–60] where by design we want to compare (and align the nodes of) equal-sized graphs. E.g., in computer vision, one might want to establish a correspondence among the nodes of two graphs, each representing a geometrical relation among m special points in two images of objects of the same type. When poses of objects do not differ significantly, the same number, m , of special points will be extracted from each image, and hence the graphs being compared will have the same size.

We can nevertheless extend our approach to graphs of different sizes. We can do so by extending two graphs, G_A and G_B , with dummy nodes such that the new graphs G'_A and G'_B have the same number of nodes. Many papers follow this approach, e.g. [61–69]. If G_A has n_A nodes and G_B has n_B nodes we can, for example, add n_B dummy nodes to G_A and n_A dummy nodes to G_B . Once we have G'_A and G'_B of equal size, we can use the methods we already described to compute a distance between G'_A and G'_B and return this distance as the distance between G_A and G_B .

Possible graph extensions differ in how the dummy nodes connect to existing graph nodes, how dummy nodes connect to themselves, and what kind of penalty we introduce for associating dummy nodes with existing graph nodes.

Method 1. One way of extending the graphs is to add dummy nodes and leave them isolated, i.e., with no edges to either existing nodes or other dummy nodes. Although this might work when both graphs are dense, it might lead to non-desirable results when one of the graphs is sparse. For example, let G_A be 3 isolated nodes and G_B be the complete graph on 4 nodes minus the edges forming triangle $\{(1,2),(2,3),(3,1)\}$. Let us assume that $S = \mathbb{P}^n$, such that, when we compute the distance between G_A and G_B , we produce an alignment between the graphs. One desirable outcome would be for G_A to be aligned with the three nodes in G_B that have no edges among them. This is basically solving the problem of finding a sparse subgraph inside a dense graph. However, computing $d_S(A', B')$, where A' and B' are the extended adjacency matrices, could equally well align G_A with the 3 dummy node of G'_B .

Method 2. Alternatively, one could add dummy nodes and connect each dummy node to all existing nodes and all other dummy nodes. This avoids the issue described for method 1. However, this creates a similar non-desirable situation: since the dummy nodes in each extended graph form a clique, we might align G_A , or G_B , with just dummy nodes, instead of producing an alignment between existing nodes in G_A and existing nodes in G_B .

Method 3. If both G_A and G_B are unweighted graphs, a method that avoids both issues above (aligning a sparse graph with isolated dummy nodes or aligning a dense graphs with cliques of dummy nodes) is to connect each dummy node to all existing nodes and all other dummy nodes with edges of weight $1/2$. This method works because, when $S = \mathbb{P}^n$, it discourages alignments of edges between existing nodes in G_A to dummy-dummy edges or dummy-existing node edges in G_B , and vice versa.

Method 4. One can also discourage aligning existing node with dummy nodes by introducing a soft linear term as in (17), penalizing mappings between dummy and existing nodes.

Method 5. Finally, a method of ensuring that the graphs have equal size is repeating them, i.e., creating “super” graphs that consist of multiple replicas of the same graph as connected components, resulting in two graphs of size equal to the least common multiple (LCM) of the sizes of the two original graphs. This is of particular use when a spectral approach, like the ones used to optimize over \mathbb{O}^n : this is because repetition, in effect, only changes the multiplicity of each value in the spectrum, which can be done (a) without affecting the spectrum structure, and (b) efficiently, once the LCM is computed.

6 Experiments

We experimentally study the properties of different graph distance measures, including metrics from our family, over several graph classes. Our main observation is that computing a heuristic estimate \hat{P} of $P^* = \arg \min_{P \in \mathbb{P}^n} \|AP - PB\|$, and using \hat{P} to estimate $d_{\mathbb{P}^n}(A, B)$ leads to violations of the metric property. In contrast, our

proposed approach of computing $d_S(A, B)$ for some S for which d a metric, and for which its computation is tractable, yields significantly improved performance in tasks such as clustering graphs (see Fig. 1).

6.1 Experimental Setup

Graphs. We use *synthetic graphs* from six classes summarized in Table 3: Barabasi Albert with degree d (Bd), Erdos Renyi with probability p (Ep), Power Law Tree (P), Regular with degree d (Rd), Small World (S), Watts Strogatz with degree d (Ws). In addition, we use a dataset of *small graphs*, comprising all 853 connected graphs of 7 nodes [70]. Finally, we use a *collaboration graph* with 5242 nodes and 14496 edges representing author collaborations [71].

Algorithms. We compare our metrics to several competitors outlined in Table 2. All receive only two unlabeled undirected simple graphs A and B and output a matching a matrix \hat{P} either in \mathbb{W}^n or in \mathbb{P}^n estimating P^* . If $\hat{P} \in \mathbb{P}^n$, we compute $\|A\hat{P} - \hat{P}B\|_1$. If $\hat{P} \in \mathbb{W}^n$, then we compute both $\|A\hat{P} - \hat{P}B\|_1$ and $\|A\hat{P} - \hat{P}B\|_F$; all norms are entry-wise. We also implement our two relaxations $d_{\mathbb{W}}$ and $d_{\mathbb{O}^n}$, for two different matrix norm combinations.

We briefly review here additional implementation details about the algorithms summarized in Table 2.

- **NetAlignBP**, **IsoRank**, **SparseIsoRank** and **NetAlignMR**, for which code is publicly available [72], are described by [14]. **Natalie** is described in [16]; code is again available [73]. All five algorithms output $P \in \mathbb{P}^n$.
- The algorithm in [15] outputs one $P \in \mathbb{P}^n$ and one $P' \in \mathbb{W}^n$. We use $P \in \mathbb{P}^n$ to compute $\|AP - PB\|_1$ and call this **InnerPerm**. We use $P' \in \mathbb{W}^n$ to compute $\|AP' - P'B\|_1$ and $\|AP' - P'B\|_2$ and call these algorithms **InnerDSL1** and **InnerDSL2** respectively. We use our own CVX-based projected gradient descent solver for the non-convex optimization problem the authors propose.
- **DSL1** and **DSL2** denote $d_S(A, B)$ when $S \in \mathbb{W}^n$ and $\|\cdot\|$ is $\|\cdot\|_1$ (element-wise) and $\|\cdot\|_F$, respectively. We implement them in Matlab (using CVX) as well as in C, aimed for medium size graphs and multi-core use. We also implemented a distributed version in Apache Spark [19] that scales to very large graphs over multiple machines based on the Alternating Directions Method of Multipliers [18].
- **ORTHOP** and **ORTHFR** denote $d_S(A, B)$ when $S \in \mathbb{O}^n$ and $\|\cdot\|$ is $\|\cdot\|_2$ (operator norm) and $\|\cdot\|_F$ respectively. We compute them using an eigendecomposition (See Sec. 4.3).
- For small graphs, we compute $d_{\mathbb{P}^n}(A, B)$ using our brute-force GPU-based code, available at [74]. For a single pair of graphs with $n \geq 15$ nodes, **EXACT** already takes several days to finish. For $\|\cdot\| = \|\cdot\|_1$ in d_S (element-wise or matrix norm), we have implemented the chemical distance as an integer value LP and solved it using branch-and-cut. It did not scale well for $n \geq 15$.
- We implemented the WL algorithm over Spark to run, multithreaded, on a machine with 40 CPUs.

We use all public algorithms as black boxes with their default parameters, as provided by the authors.

6.2 Experimental Results

Clustering Graphs. The difference between our metrics and non-metrics is striking when clustering graphs. This is illustrated by the clustering experiment shown in Fig. 1. Graphs of size $n = 50$ from the 6 classes in Table 3 are clustered together through hierarchical agglomerative clustering. We compute distances between them using nine different algorithms; only the distances in our family (DSL1, DSL2, ORTHOP, and ORTHFR) are metrics. The quality of clusters induced by our metrics are far superior than clusters induced by non-metrics; in fact, **ORTHOP** and **ORTHFR** can lead to no misclassifications. This experiment strongly suggests our produced metrics correctly capture the topology of the metric space between these larger graphs.

Triangle Inequality Violations. Given graphs A, B and C and a distance d , a Triangle Inequality Violation (TIV) occurs when $d(A, C) > d(A, B) + d(B, C)$. Being metrics, none of our distances induce TIVs; this is not the case for the remaining algorithms in Table 2. Fig. 2 shows the TIV fraction across the synthetic graphs of Table 3 while Fig. 3 shows the fraction of TIVs found on the 853 small graphs ($n = 7$). **NetAlignMR** also produces no TIVs on the small graphs, but it does induce TIVs in synthetic graphs. We observe that it is easier to find TIVs when graphs are close: in synthetic graphs, TIVs abound for $n = 10$. No algorithm performs well across all categories of graphs.

Effect of TIVs on Clustering. Next, to investigate the effect of TIVs on clustering, we artificially introduced triangle inequality violations into the pairs of distances between graphs. We then re-evaluated clustering performance for hierarchical agglomerative clustering using the *Ward* method, which performed best in Fig. 1. Fig. 4 shows the fraction of misclassified graphs as the fraction of TIVs introduced increases. To incur as small a perturbation on distances as possible, we introduce TIVs as follows: For every three graphs, A, B, C , with probability p , we set $d(A, C) = d(A, B) + d(B, C)$. Although this does not introduce a TIV w.r.t. A, B , and C , this distortion does introduce TIVs w.r.t. other triplets involving A and C . We repeat this 20 times for each algorithm and each value of p , and compute the average fraction of TIVs, shown in the x -axis, and the average fraction of misclassified graphs, shown in the y -axis. As little as 1% TIVs significantly deteriorate clustering performance. Note that the fraction of TIVs is computed over the total number of TIVs possible, which grows cubically with the number of graphs being clustered. We also see that, even after introducing TIVs, clustering based on metrics outperforms clustering based on non-metrics.

Comparison to Chemical Distance. We compare how different distance scores relate to the chemical distance **EXACT** through two experiments on the small graphs (computation on larger graphs is prohibitive). In Figure 5(a), we compare the distances between small graphs with 7 nodes produced by the different algorithms and **EXACT** using the DISTATIS method of [75]. Let $D \in \mathbb{R}_+^{835 \times 835}$ be the matrix of distances between graphs under an algorithm. DISTATIS computes the normalized Laplacian of this matrix, given by $L = -UDU/\|UDU\|_2$ where $U = I - \frac{\mathbf{1}\mathbf{1}^T}{n}$.

The DISTATIS score is the cosine similarity of such Laplacians (vectorized). We see that our metrics produce distances attaining high similarity with **EXACT**, though **NetAlignBP** has the highest similarity. We measure proximity to **EXACT** with an additional test. Given D , we compute the nearest neighbor (NN) meta-graph by connecting a graph in D to every graph at distance less than its average distance to other graphs. This results in a (labeled) meta-graph, which we can compare to the NN meta-graph induced by other algorithms, measuring the fraction of distinct edges. Fig. 5(b) shows that our algorithms perform quite well, though **Natalie** yields the smallest distance to **EXACT**.

Incorporating Constraints. Computation costs can be reduced through metric embeddings, as in (17). To show this, we produce a copy of the 5242 node collaboration graph with permuted node labels. We then run the WL algorithm [47] to produce structural colors, which induce coloring constraints on $P \in \mathbb{W}^n$. The WL algorithm reaches a fixed point after $k = 5$ iterations. The support of P (i.e., the number of variables in the optimization (12)), the support of $AP - PA$ (i.e., the number of non-zero summation terms in the objective of (12)), as well as the execution time τ of the WL algorithm, are summarized in Table 4. The original unconstrained problem involves $5242^2 \approx 27.4\text{M}$ variables. However, after using WL and induced constraints, the effective dimension of the optimization problem (12) reduces considerably. This, in turn, speeds up convergence time, shown in Fig. 6: including the time to compute constraints, a solution is found 110 times faster after the introduction of the constraints.

7 Conclusion

Our work suggests that incorporating soft and hard constraints has a great potential to further improve the efficiency of our metrics. In future work, we intend to investigate and characterize the resulting equivalence classes under different soft and hard constraints, and to quantify these gains in efficiency. We also plan to develop scalable distributed solvers for our family of metrics. A good starting point is the Alternating Direction Method of Multipliers [76, 77], which enjoys several useful properties. Specifically, under proper tuning and mild convexity assumptions, it achieves the convergence rate of the fastest-possible first-order method [78, 79], it can be less affected by the topology of the communication network in a cluster than, e.g. gradient descent [80, 81], and it parallelizes well both on share-memory multiprocessor systems, GPUs and computer clusters [18, 82, 83]. Determining the necessity of the conditions used in proving that d_S is a metric is also an open problem. Finally, we are investigating generalizations of our family of metrics to multi-metrics, i.e. we want to define a tractable closeness score for a set of $n > 2$ graphs that satisfies a generalization of the properties of metrics for more than two elements [84].

Abbreviations

ADMM: Alternating Directions Method of Multipliers, CKS: Chartrand-Kubiki-Shultz, GPU: Graph Processing Unit, LCM:Least Common Multiple, WL: Weisfeiler-Lehman

Declarations

Acknowledgements

The authors gratefully acknowledge the support of the National Science Foundation (grants IIS-1741197, IIS-1741129) and of the National Institutes of Health (grant 1U01AI124302).

Availability of data and materials

McKay, B.: List of 7 node connected graphs. <http://users.cecs.anu.edu.au/~bdm/data/graphs.html>
 Leskovec, J., Kleinberg, J., Faloutsos, C.: Stanford Large Network Dataset Collection. <http://snap.stanford.edu/data/ca-GrQc.html>
 Khan, A., Gleich, D., Halappanavar, M., Pothen, A.: Multicore codes for network alignment. <https://www.cs.purdue.edu/homes/dgleich/codes/netalignmc/>
 El-Kebir, M., Heringa, J., Klau, G.: Natalie, a tool for pairwise global network alignment. <http://www.mi.fu-berlin.de/w/LISA/Natalie>
 Bento, J.: Graph distance via brute force GPU computation. https://github.com/bentoayr/exact_graph_match

Competing Interests

The authors declare that they have no competing interests.

Author Contributions

The authors proved the main theorems collaboratively. All experimental results were implemented and executed by Professor Bento, with the exception of the experiments in Fig. 6, which was implemented and executed by Professor Ioannidis.

Author details

¹Department of Computer Science, Boston College, St. Mary's Hall, 2nd floor, 02467 Chestnut Hill, MA, USA.

²Department of Electrical and Computer Engineering, Northeastern University, 360 Huntington Avenue, 02115 Boston, MA, USA.

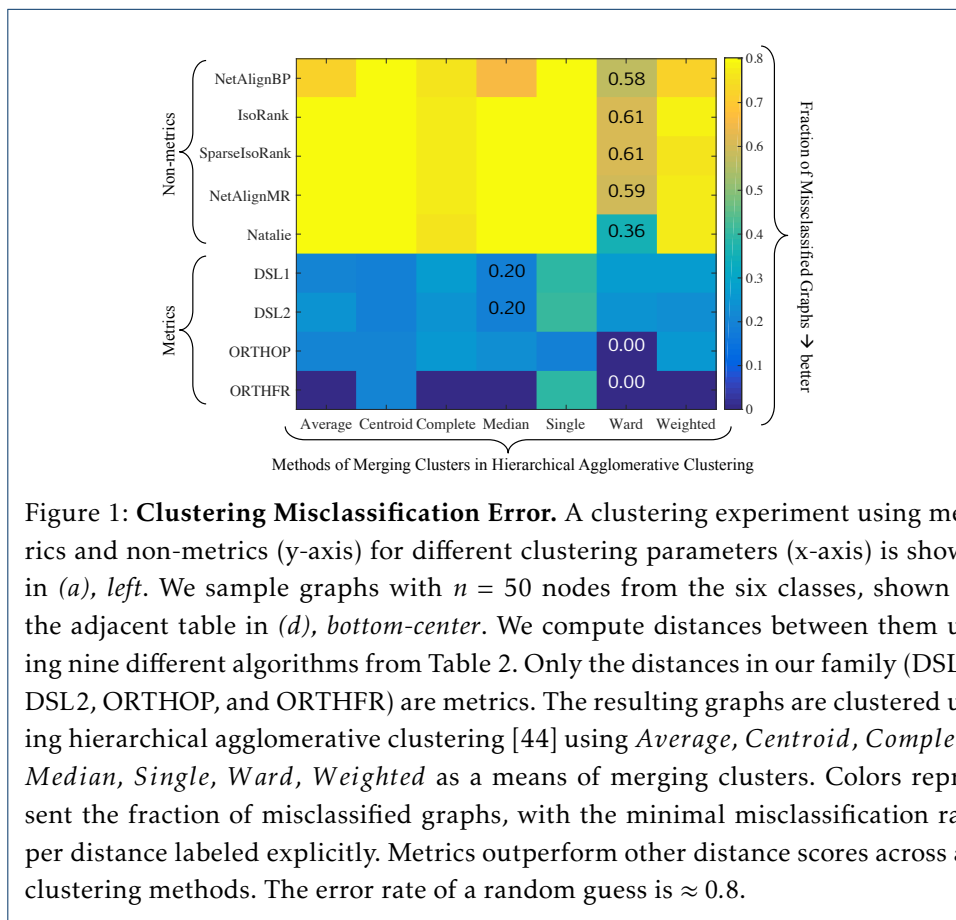
References

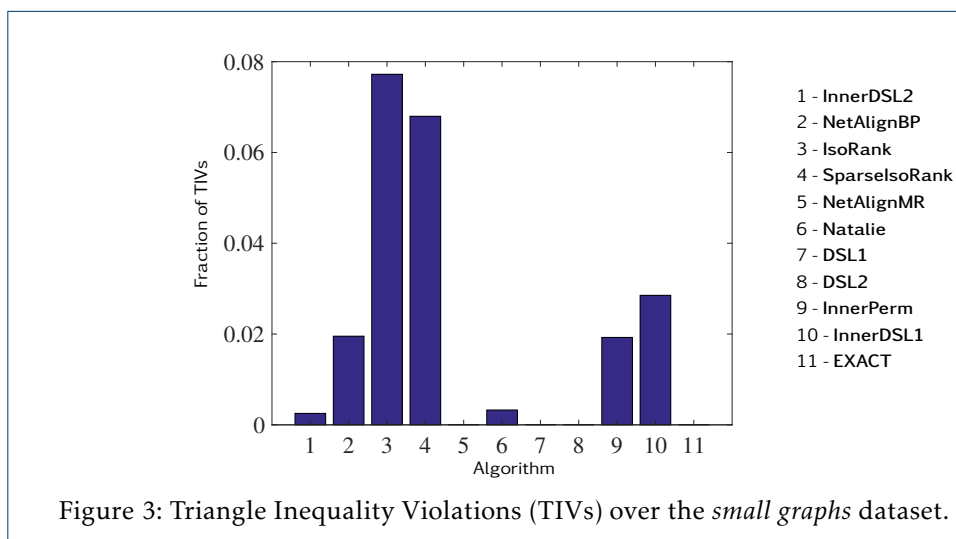
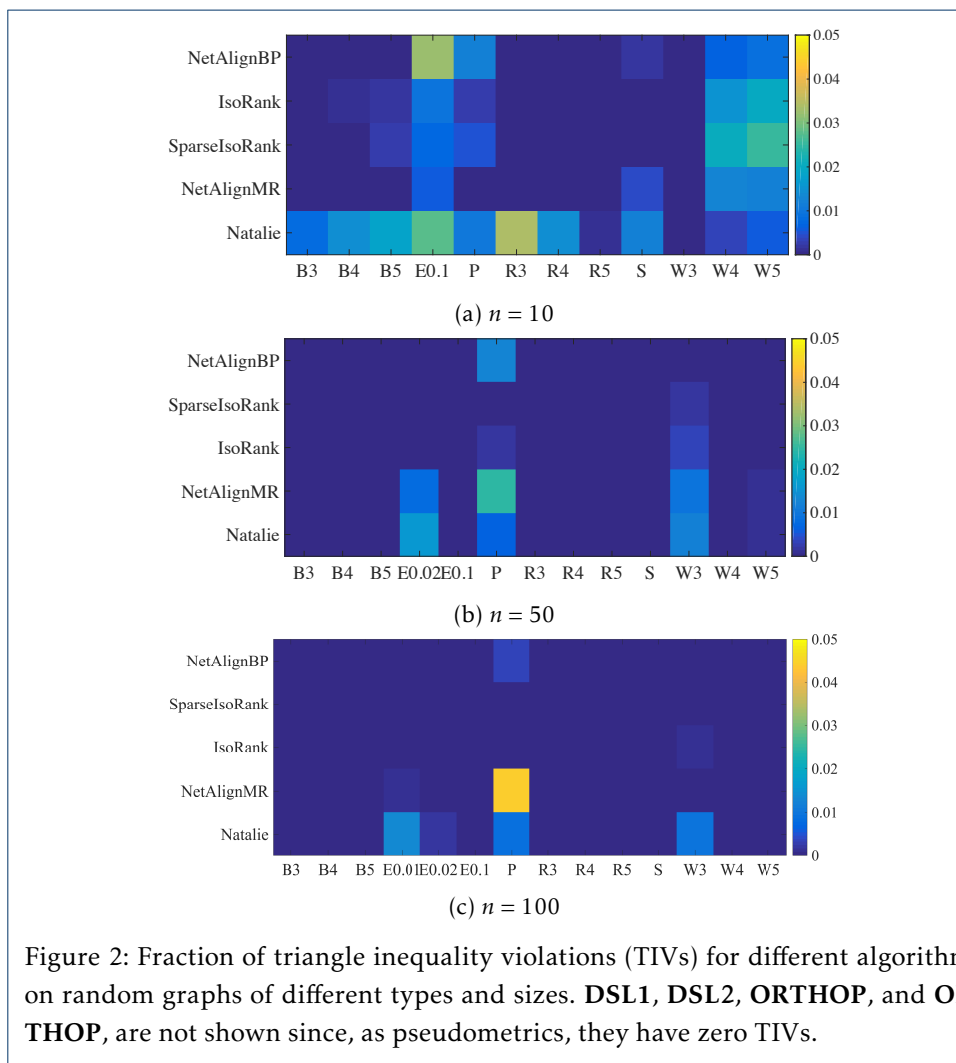
- Conte, D., Foggia, P., Sansone, C., Vento, M.: Thirty years of graph matching in pattern recognition. *International Journal of Pattern Recognition and Artificial Intelligence* **18**(03), 265–298 (2004)
- Macindoe, O., Richards, W.: Graph comparison using fine structure analysis. In: *SocialCom* (2010)
- Koutra, D., Vogelstein, J.T., Faloutsos, C.: Deltacon: A principled massive-graph similarity function. In: *SDM* (2013)
- Clarkson, K.L.: Nearest-neighbor searching and metric space dimensions. *Nearest-Neighbor Methods for Learning and Vision: Theory and Practice*, 15–59 (2006)
- Clarkson, K.L.: Nearest neighbor queries in metric spaces. *Discrete & Computational Geometry* **22**(1), 63–93 (1999)
- Beygelzimer, A., Kakade, S., Langford, J.: Cover trees for nearest neighbor. In: *ICML* (2006)
- Ackermann, M.R., Blömer, J., Sohler, C.: Clustering for metric and nonmetric distance measures. *ACM Transactions on Algorithms (TALG)* **6**(4), 59 (2010)
- Angiulli, F., Pizzuti, C.: Fast outlier detection in high dimensional spaces. In: *PKDD* (2002)
- Indyk, P.: Sublinear time algorithms for metric space problems. In: *Proceedings of the Thirty-first Annual ACM Symposium on Theory of Computing*, pp. 428–434 (1999). *ACM*
- Kvasnička, V., Pospíchal, J., Baláz, V.: Reaction and chemical distances and reaction graphs. *Theoretical Chemistry Accounts: Theory, Computation, and Modeling (Theoretica Chimica Acta)* **79**(1), 65–79 (1991)
- Chartrand, G., Kubicki, G., Schultz, M.: Graph similarity and distance in graphs. *Aequationes Mathematicae* **55**(1–2), 129–145 (1998)
- Babai, L.: Graph isomorphism in quasipolynomial time [extended abstract]. In: *STOC* (2016)
- Klau, G.W.: A new graph-based method for pairwise global network alignment. *BMC bioinformatics* **10**(1), 59 (2009)
- Bayati, M., Gerritsen, M., Gleich, D.F., Saberi, A., Wang, Y.: Algorithms for large, sparse network alignment problems. In: *ICDM* (2009)
- Lyzinski, V., Fishkind, D.E., Fiori, M., Vogelstein, J.T., Priebe, C.E., Sapiro, G.: Graph matching: Relax at your own risk. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **38**(1), 60–73 (2016)
- El-Kebir, M., Heringa, J., Klau, G.W.: Natalie 2.0: Sparse global network alignment as a special case of quadratic assignment. *Algorithms* **8**(4), 1035–1051 (2015)
- Umeyama, S.: An eigendecomposition approach to weighted graph matching problems. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **10**(5), 695–703 (1988)
- Boyd, S., Parikh, N., Chu, E., Peleato, B., Eckstein, J.: Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine Learning* **3**(1), 1–122 (2011)
- Zaharia, M., Chowdhury, M., Franklin, M.J., Shenker, S., Stoica, I.: Spark: Cluster computing with working sets. *HotCloud* **10**(10–10), 95 (2010)
- Allen, F.H.: The Cambridge Structural Database: a quarter of a million crystal structures and rising. *Acta Crystallographica Section B: Structural Science* **58**(3), 380–388 (2002)
- Papadimitriou, P., Dasdan, A., Garcia-Molina, H.: Web graph similarity for anomaly detection. *Journal of Internet Services and Applications* **1**(1), 19–30 (2010)
- Soundarajan, S., Eliassi-Rad, T., Gallagher, B.: A guide to selecting a network similarity method. In: *SDM* (2014)
- Garey, M.R., Johnson, D.S.: *Computers and Intractability* vol. 29. W. H. Freeman and Company, New York, USA (2002)
- Sanfeliu, A., Fu, K.: A distance measure between attributed relational graphs for pattern recognition. *IEEE Transactions on Systems, Man, and Cybernetics* (3), 353–362 (1983)

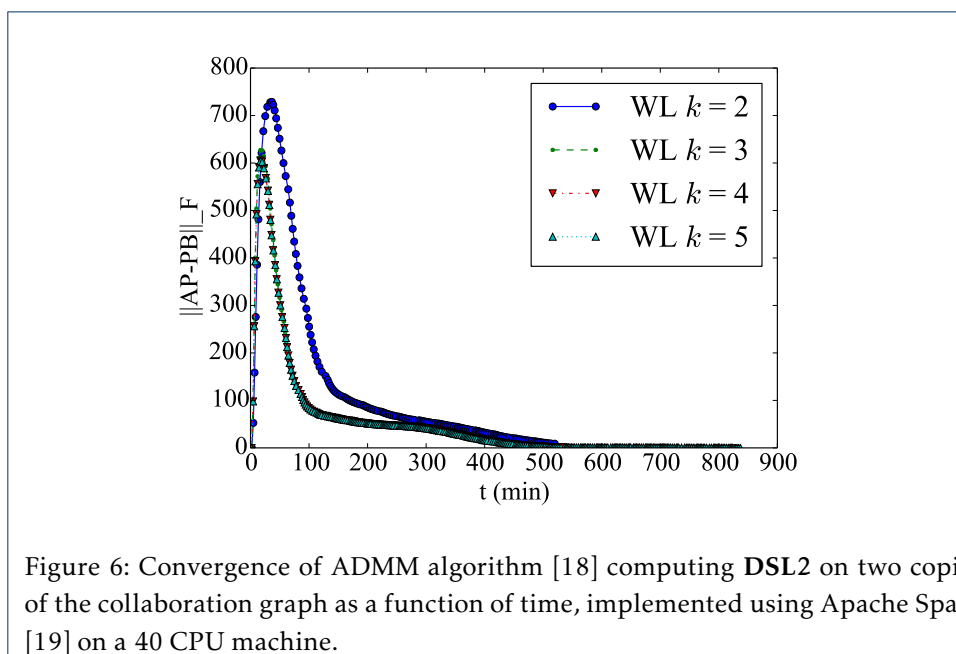
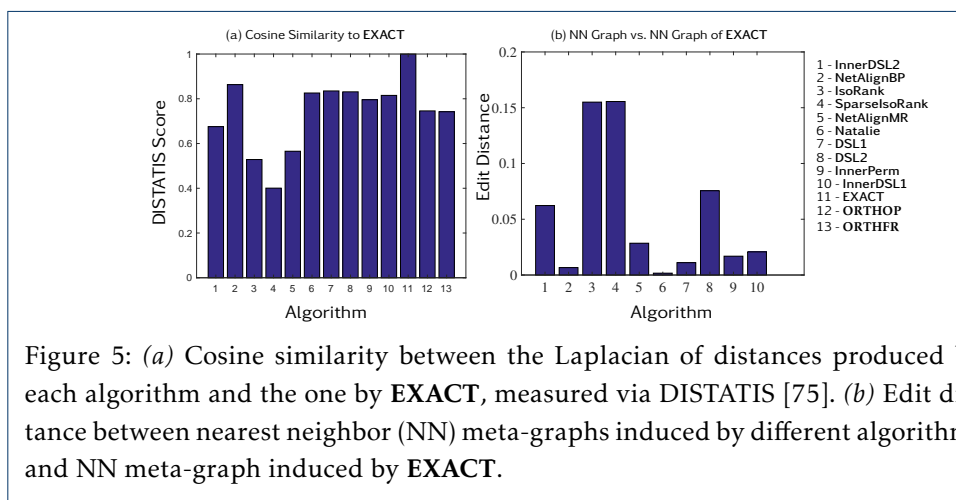
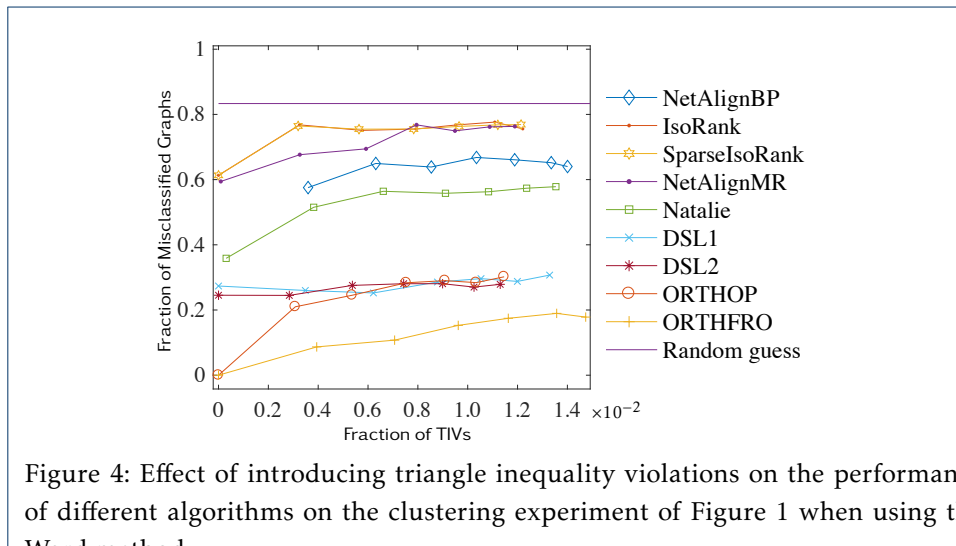
25. Bunke, H., Shearer, K.: A graph distance metric based on the maximal common subgraph. *Pattern Recognition Letters* **19**(3), 255–259 (1998)
26. Bunke, H.: On a relation between graph edit distance and maximum common subgraph. *Pattern Recognition Letters* **18**(8), 689–694 (1997)
27. Riesen, K., Bunke, H.: Approximate graph edit distance computation by means of bipartite graph matching. *Image and Vision Computing* **27**(7), 950–959 (2009)
28. Fankhauser, S., Riesen, K., Bunke, H.: Speeding up graph edit distance computation through fast bipartite matching. In: *GBR* (2011)
29. Koca, J., Kratochvil, M., Kvasnicka, V., Matyska, L., Pospichal, J.: *Synthon Model of Organic Chemistry and Synthesis Design* vol. 51. Springer, Berlin/Heidelberg, Germany (2012)
30. Jain, B.J.: On the geometry of graph spaces. *Discrete Applied Mathematics* **214**, 126–144 (2016)
31. Riesen, K., Neuhaus, M., Bunke, H.: Graph embedding in vector spaces by means of prototype selection. In: *GBR* (2007)
32. Ferrer, M., Valveny, E., Serratos, F., Riesen, K., Bunke, H.: Generalized median graph computation by means of graph embedding in vector spaces. *Pattern Recognition* **43**(4), 1642–1655 (2010)
33. Riesen, K., Bunke, H.: *Graph Classification and Clustering Based on Vector Space Embedding* vol. 77. World Scientific, Singapore (2010)
34. Zhu, P., Wilson, R.C.: A study of graph spectra for comparing graphs. In: *BMVC* (2005)
35. Wilson, R.C., Zhu, P.: A study of graph spectra for comparing graphs and trees. *Pattern Recognition* **41**(9), 2833–2841 (2008)
36. Elghawalby, H., Hancock, E.R.: Measuring graph similarity using spectral geometry. In: *ICIAR* (2008)
37. Aflalo, Y., Bronstein, A., Kimmel, R.: On convex relaxation of graph isomorphism. *PNAS* **112**(10), 2942–2947 (2015)
38. Koutra, D., Tong, H., Lubensky, D.: Big-align: Fast bipartite graph alignment. In: *ICDM* (2013)
39. Schellewald, C., Roth, S., Schnörr, C.: Evaluation of convex optimization techniques for the weighted graph-matching problem in computer vision. In: *Pattern Recognition* (2001)
40. Sebastian, T.B., Klein, P.N., Kimia, B.B.: Recognition of shapes by editing their shock graphs. *IEEE Transactions on pattern analysis and machine intelligence* **26**(5), 550–571 (2004)
41. He, L., Han, C.Y., Wee, W.G.: Object recognition and recovery by skeleton graph matching. In: *ICME* (2006)
42. Vogelstein, J.T., Conroy, J.M., Podrazik, L.J., Kratzer, S.G., Harley, E.T., Fishkind, D.E., Vogelstein, R.J., Priebe, C.E.: Large (brain) graph matching via fast approximate quadratic programming. *arXiv preprint arXiv:1112.5507* (2011)
43. Xing, E.P., Ng, A.Y., Jordan, M.I., Russell, S.: Distance metric learning with application to clustering with side-information. In: *NIPS*, vol. 15, p. 12 (2002)
44. Hartigan, J.A.: *Clustering Algorithms*. Wiley, New York (1975)
45. Bento, J., Ioannidis, S.: A family of tractable graph distances. In: *Proceedings of the 2018 SIAM International Conference on Data Mining*, pp. 333–341 (2018). SIAM
46. Birkhoff, G.: Three observations on linear algebra. *Univ. Nac. Tucumán. Revista A* **5**, 147–151 (1946)
47. Weisfeiler, B., Lehman, A.A.: A reduction of a graph to a canonical form and an algebra arising during this reduction. *Nauchno-Tekhnicheskaya Informatsia* **2**(9), 12–16 (1968)
48. Boyd, S., Vandenberghe, L.: *Convex Optimization*. Cambridge university press, Cambridge, UK (2004)
49. Nesterov, Y., Nemirovskii, A.: *Interior-point Polynomial Algorithms in Convex Programming* vol. 13. Siam, ??? (1994)
50. Bertsekas, D.P.: Nonlinear programming. *Journal of the Operational Research Society* **48**(3), 334–334 (1997)
51. Absil, P.-A., Mahony, R., Sepulchre, R.: *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, Princeton, NJ, USA (2009)
52. Ramana, M.V., Scheinerman, E.R., Ullman, D.: Fractional isomorphism of graphs. *Discrete Mathematics* **132**(1-3), 247–265 (1994)
53. Tinhofer, G.: Graph isomorphism and theorems of Birkhoff type. *Computing* **36**(4), 285–300 (1986)
54. Kuhn, H.W.: The hungarian method for the assignment problem. *Naval Research Logistics Quarterly* **2**(1-2), 83–97 (1955)
55. Jonker, R., Volgenant, A.: A shortest augmenting path algorithm for dense and sparse linear assignment problems. *Computing* **38**(4), 325–340 (1987)
56. Hoffman, A.J., Wielandt, H.W.: The variation of the spectrum of a normal matrix. *Duke Math. J* **20**(1), 37–39 (1953)
57. Horn, R.A., Johnson, C.R.: *Matrix Analysis*. Cambridge University Press, Cambridge, UK (2012)
58. Hu, N., Thibert, B., Guibas, L.: Distributable consistent multi-graph matching. *arXiv preprint arXiv:1611.07191* (2016)
59. Shen, Y., Lin, W., Yan, J., Xu, M., Wu, J., Wang, J.: Person re-identification with correspondence structure learning. In: *Proceedings of the IEEE International Conference on Computer Vision*, pp. 3200–3208 (2015)
60. Pachauri, D., Kondor, R., Singh, V.: Solving the multi-way matching problem by permutation synchronization. In: *Advances in Neural Information Processing Systems*, pp. 1860–1868 (2013)
61. Zaslavskiy, M., Bach, F., Vert, J.P.: Global alignment of protein–protein interaction networks by graph matching methods. *Bioinformatics* **25**(12), 259–267 (2009)
62. Zaslavskiy, M., Bach, F., Vert, J.-P.: A path following algorithm for the graph matching problem. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **31**(12), 2227–2242 (2009)
63. Narayanan, A., Shi, E., Rubinstein, B.I.: Link prediction by de-anonymization: How we won the kaggle social network challenge. In: *Neural Networks (IJCNN), The 2011 International Joint Conference On*, pp. 1825–1834 (2011). IEEE
64. Zaslavskiy, M., Bach, F., Vert, J.-P.: Many-to-many graph matching: a continuous relaxation approach. In: *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pp. 515–530 (2010). Springer

65. Zhou, F., De la Torre, F.: Factorized graph matching. In: IEEE Conference on Computer Vision & Pattern Recognition (CVPR), pp. 127–134 (2012). IEEE
66. Gold, S., Rangarajan, A.: Softmax to softassign: Neural network algorithms for combinatorial optimization. *Journal of Artificial Neural Networks* 2(4), 381–399 (1996)
67. Yan, J., Cho, M., Zha, H., Yang, X., Chu, S.: A general multi-graph matching approach via graduated consistency-regularized boosting. arXiv preprint arXiv:1502.05840 (2015)
68. Solé-Ribalta, A., Serratos, F.: Graduated assignment algorithm for finding the common labelling of a set of graphs. In: Joint IAPR International Workshops on Statistical Techniques in Pattern Recognition (SPR) and Structural and Syntactic Pattern Recognition (SSPR), pp. 180–190 (2010). Springer
69. Yan, J., Wang, J., Zha, H., Yang, X., Chu, S.: Consistency-driven alternating optimization for multigraph matching: A unified approach. *IEEE Transactions on Image Processing* 24(3), 994–1009 (2015)
70. McKay, B.: List of 7 node connected graphs. <http://users.cecs.anu.edu.au/~bdm/data/graphs.html>
71. Leskovec, J., Kleinberg, J., Faloutsos, C.: Stanford Large Network Dataset Collection. <http://snap.stanford.edu/data/ca-GrQc.html>
72. Khan, A., Gleich, D., Halappanavar, M., Pothan, A.: Multicore codes for network alignment. <https://www.cs.purdue.edu/homes/dgleich/codes/netalignmc/>
73. El-Kebir, M., Heringa, J., Klau, G.: Natalie, a tool for pairwise global network alignment. <http://www.mi.fu-berlin.de/w/LISA/Natalie>
74. Bento, J., Ioannidis, S.: Graph distance via brute force GPU computation. https://github.com/bentoayr/exact_graph_match
75. Abdi, H., O’Toole, A.J., Valentin, D., Edelman, B.: DISTATIS: The analysis of multiple distance matrices. In: CVPR Workshops (2005)
76. Gabay, D., Mercier, B.: A dual algorithm for the solution of nonlinear variational problems via finite element approximation. *Computers & Mathematics with Applications* 2(1), 17–40 (1976)
77. Glowinski, R., Marroco, A.: Sur l’approximation, par éléments finis d’ordre un, et la résolution, par pénalisation-dualité d’une classe de problèmes de dirichlet non linéaires. *ESAIM: Mathematical Modelling and Numerical Analysis-Modélisation Mathématique et Analyse Numérique* 9(R2), 41–76 (1975)
78. França, G., Bento, J.: An explicit rate bound for over-relaxed admm. In: Information Theory (ISIT), 2016 IEEE International Symposium On, pp. 2104–2108 (2016). IEEE
79. Nesterov, Y.: *Introductory Lectures on Convex Optimization: A Basic Course* vol. 87. Springer, Berlin/Heidelberg, Germany (2013)
80. França, G., Bento, J.: Markov chain lifting and distributed ADMM. *IEEE Signal Processing Letters* 24(3), 294–298 (2017)
81. França, G., Bento, J.: How is distributed admm affected by network topology? arXiv preprint arXiv:1710.00889 (2017)
82. Parikh, N., Boyd, S.: Block splitting for distributed optimization. *Mathematical Programming Computation* 6(1), 77–102 (2014)
83. Hao, N., Oghbaee, A., Rostami, M., Derbinsky, N., Bento, J.: Testing fine-grained parallelism for the admm on a factor-graph. In: Parallel and Distributed Processing Symposium Workshops, 2016 IEEE International, pp. 835–844 (2016). IEEE
84. Safavi, S., Bento, J.: n-metrics for multiple graph alignment. arXiv preprint arXiv:1807.03368 (2018)
85. Singh, R., Xu, J., Berger, B.: Pairwise global alignment of protein interaction networks by matching neighborhood topology. In: RECOMB (2007)
86. Albert, R., Barabási, A.-L.: Statistical mechanics of complex networks. *Reviews of Modern Physics* 74(1), 47 (2002)
87. Erdős, P., Rényi, A.: On random graphs, i. *Publicationes Mathematicae (Debrecen)* 6, 290–297 (1959)
88. Mahmoud, H.M., Smythe, R.T., Szymański, J.: On the structure of random plane-oriented recursive trees and their branches. *Random Structures & Algorithms* 4(2), 151–176 (1993)
89. Bollobás, B.: Random graphs. In: *Modern Graph Theory*, pp. 215–252. Springer, Berlin/Heidelberg, Germany (1998)
90. Kleinberg, J.: The small-world phenomenon: An algorithmic perspective. In: STOC (2000)
91. Watts, D.J., Strogatz, S.H.: Collective dynamics of ‘small-world’ networks. *Nature* 393(6684), 440–442 (1998)

Figures







Tables

Table 1: Notation Summary

$[n]$	Set $\{1, \dots, n\}$
$\mathbb{R}^{n \times n}$	The set of real $n \times n$ matrices.
\mathbf{S}^n	The set of real, symmetric matrices.
I	The identity matrix of size $n \times n$.
$\mathbf{1}$	The n -dimensional vector whose entries are all equal to 1.
$\sigma_{\max}(\cdot)$	Largest singular value of a matrix.
$\text{tr}(\cdot)$	The trace of a matrix.
$\text{conv}(\cdot)$	The convex hull of a set.
$G(V, E)$	Graph with vertex set V and edge set E .
A, B	Matrices $[a_{i,j}]_{i,j \in [n]}, [b_{i,j}]_{i,j \in [n]}$.
$\ \cdot\ _p$	Operator or entry-wise p -norm.
$\ \cdot\ _F$	Frobenius norm.
\mathbb{P}^n	Set of permutation matrices of size $n \times n$, c.f. (4)
\mathbb{W}^n	Set of doubly stochastic matrices (a.k.a. the Birkhoff polytope) of size $n \times n$, c.f. (5)
\mathbb{O}^n	Set of orthonormal matrices (a.k.a. the Stiefel manifold) of size $n \times n$, c.f. (6)
$\Omega, \tilde{\Omega}$	Sets over which a metric is defined.
$d(x, y)$	A metric over space Ω .
$\tilde{d}(x, y)$	The symmetric extension of $d(x, y)$.
(Ω, d)	A metric space.
G_A, G_B	Graphs with adjacency matrices A, B .
P, W, O	$n \times n$ matrices.
S	A closed and bounded subset of $\mathbb{R}^{n \times n}$.
$d_S(A, B)$	A class of distance scores defined by minimization (12) over set S .
$d_{\mathbb{P}^n}$	Pseudometric d_S , where S is the set of permutation matrices.
$d_{\mathbb{W}^n}$	Pseudometric d_S , where S is the set of doubly stochastic matrices.
$d_{\mathbb{O}^n}$	Pseudometric d_S , where S is the set of orthogonal matrices.
$\Psi_{\tilde{\Omega}}^n$	Set of all embeddings from $[n] \rightarrow \tilde{\Omega}$, where $(\tilde{\Omega}, \tilde{d})$ is a metric space.
ψ_A, ψ_B	Embeddings in $\Psi_{\tilde{\Omega}}^n$ of nodes in graphs G_A and G_B , respectively.
D_{ψ_A, ψ_B}	$n \times n$ matrix of all pairwise distances between images of nodes in G_A and G_B , under embeddings ψ_A and ψ_B .

Table 2: Competitor Distance Scores & Our Metrics

(Non-metric) Distance Score Algorithms	
NetAlignBP	Network Alignment using Belief Propagation [14, 72]
IsoRank	Neighborhood Topology Isomorphism using Page Rank [72, 85]
SparselsoRank	Neighborhood Topology Sparse Isomorphism using Page Rank [14, 72]
InnerPerm	Inner Product Matching with Permutations [15]
InnerDSL1	Inner Product Matching with Matrices in \mathbb{W}^n and entry-wise 1-norm [15]
InnerDSL2	Inner Product Matching with Matrices in \mathbb{W}^n and Frobenius norm [15]
NetAlignMR	Iterative Matching Relaxation [13, 72]
Natalie (V2.0)	Improved Iterative Matching Relaxation [16, 73]
Metrics from our Family (2)	
EXACT	Chemical Distance via brute force search over GPU
DSL1	Doubly Stochastic Chemical Distance $d_{\mathbb{W}^n}$ with entry-wise 1-norm
DSL2	Doubly Stochastic Chemical Distance $d_{\mathbb{W}^n}$ with Frobenius norm
ORTHOP	Orthogonal Relaxation of Chemical Distance $d_{\mathbb{O}^n}$ with operator 2-norm
ORTHFR	Orthogonal Relaxation of Chemical Distance $d_{\mathbb{O}^n}$ with Frobenius norm

Table 3: Synthetic Graph Classes

Description	
Bd	Barabasi Albert of degree d [86]
Ep	Erdős-Rényi with probability p [87]
P	Power Law Tree [88]
Rd	Regular Graph of degree d [89]
S	Small World [90]
Wd	Watts Strogatz of degree d [91]

Table 4: Effect of coloring/hard constraints.

k	$\ P\ _0$	$\ AP-PA\ _0$	τ
1	3,747,960	100.569	133s
2	239,048	3,004	104s
3	182,474	2,036	136s
4	182,016	2,030	169s
5	182,006	2,030	200s

Effect of coloring/hard constraints on the numbers of variables ($\|P\|_0$) and terms of objective ($\|AP-PA\|_0$) using k iterations of the WL coloring algorithm. The last column shows the execution time of WL on a 40 CPU machine using Apache Spark [19].