
Exact Shape Correspondence via 2D Graph Convolution

Barakeel Fansu Kamhoua¹, Lin Zhang^{2*}, Yongqiang Chen¹, Han Yang¹, Kaili Ma¹,
Bo Han³, Bo Li², James Cheng¹

¹The Chinese University of Hong Kong, ²The Hong Kong University of Science and Technology,
³Hong Kong Baptist University

Abstract

For exact 3D shape correspondence (matching or alignment), i.e., the task of matching each point on a shape to its exact corresponding point on the other shape (or to be more specific, matching at geodesic error 0), most existing methods do not perform well due to two main problems. First, on nearly-isometric shapes (i.e., low noise levels), most existing methods use the eigen-vectors (eigen-functions) of the Laplace Beltrami Operator (LBO) or other shape descriptors to update an initialized correspondence which is not exact, leading to an accumulation of update errors. Thus, though the final correspondence may generally be smooth, it is generally inexact. Second, on non-isometric shapes (noisy shapes), existing methods are generally not robust to noise as they usually assume near-isometry. In addition, existing methods that attempt to address the non-isometric shape problem (e.g., GRAMPA) are generally computationally expensive and do not generalise to nearly-isometric shapes. To address these two problems, we propose a 2D graph convolution-based framework called 2D-GEM. 2D-GEM is robust to noise on non-isometric shapes and with a few additional constraints, it also addresses the errors in the update on nearly-isometric shapes. We demonstrate the effectiveness of 2D-GEM by achieving a high accuracy of 90.5% at geodesic error 0 on the non-isometric benchmark SHREC16, i.e., TOPKIDS (while being much faster than GRAMPA), and on nearly-isometric benchmarks by achieving a high accuracy of 92.5% on TOSCA and 84.9% on SCAPE at geodesic error 0.

1 Introduction

3D shapes can undergo different types of deformations. Deformations can be rigid or non-rigid. Rigid deformations include rotations, translations, reflections, or any composition of these resulting in pose changes. They preserve geometric properties of the shape such as segment lengths, angles, volumes, sizes, geodesic distances and others. Non-rigid deformations do not preserve geometric properties of the shape, they include dilations, or a composition of dilations and rigid deformations (e.g., stretching, scaling). In this work we focus on non-rigid shapes that are either nearly-isometric (i.e., most geometric properties are preserved), or non-isometric (i.e., geometric properties are not preserved, and contain strong topological noise).

Non-rigid shape matching is a very important task in machine learning, shape analysis and computer vision. It has important applications such as shape registration, comparison, recognition, and retrieval. 3D shapes are usually represented as 2D manifolds embedded in \mathbb{R}^3 . Given two three dimensional shapes \mathcal{M}_1 and \mathcal{M}_2 with n vertices each, the goal of non-rigid shape matching is to find a meaningful correspondence $\mathcal{T} : \mathcal{M}_1 \rightarrow \mathcal{M}_2$, where \mathcal{T} is (a) bijective, (b) continuous in both directions in the sense that nearby points on \mathcal{M}_1 should be matched to nearby points on \mathcal{M}_2 (vice versa), and (c)

*Corresponding author: lzhangbv@connect.ust.hk

similar points should be matched to each other [42]. In the discrete setting the correspondence (map or match) \mathcal{T} can be represented as a permutation matrix $\mathbf{P} \in \{0, 1\}^{n \times n}$ [54, 4].

In more details, 3D shapes are usually represented using triangulations (e.g., the Delauney triangulation [20]) and xyz-coordinates representations in \mathbb{R}^3 [80, 42]. It is hard to utilize these xyz-coordinates and triangulations to directly perform 3D shape matching because of: (a) embedding ambiguities, i.e. shapes sharing the same metric (isometry) or very similar metrics (nearly-isometric) can have drastically different xyz-coordinates representations in \mathbb{R}^3 due to pose changes [80], and (b) non-isometric shapes with strong topological noise will have even more different xyz-coordinates and triangulations.

As such, these 3D-coordinates together with the triangulations are usually used to extract *point-wise descriptors* (which encode the similarity between points), and *pair-wise descriptors* (which encode global and/or local relations between pairs of points). Once these descriptors are built, the shape matching problem is reduced to an alignment problem over (a) the *point-wise descriptors*, or (b) the *pair-wise descriptors*, or (c) *both (joint alignment)*. The *joint alignment* problem can be written as:

$$\mathbf{P} = \arg \min_{\mathbf{P} \in \Pi} \mathbb{E}(\mathbf{P}) = \arg \min_{\mathbf{P} \in \Pi} \mathbb{F}(\mathbf{P}) + \alpha \mathbb{H}(\mathbf{P}), \quad (1)$$

where Π is the set of permutation matrices, \mathbb{E} is the generalised energy, \mathbb{F} is the energy over *point-wise descriptors*, \mathbb{H} is the energy over *pair-wise descriptors*, and α is a trade-off parameter [42]. Several algorithms have been proposed to address point-wise alignment, [4, 70, 11, 55, 64], pair-wise alignment [20, 18], as well as joint alignment [59, 57, 56, 2, 79, 42, 54, 58]. Due to space constraints, see Appendix A for more details on the different alignments as well as related works.

To measure the performance of an algorithm on the 3D shape matching task, the geodesic error is usually used as the error metric [20, 80]. Given that a point x on shape \mathcal{M}_1 is mapped by an algorithm to y as its correspondence on shape \mathcal{M}_2 , the geodesic error measures how far the match y is from the true matching position for x on \mathcal{M}_2 i.e., y^* (given by the true map). When the map returned by algorithm is exact, it significantly improves performance and efficiency of other tasks such as registration, comparison, motion transfer, style transfer. However, most methods proposed in each of the aforementioned categories for 3D shape matching do not perform well at geodesic error 0 (i.e., the task of exact matching where y is the same point y^*). This is mainly due to two factors.

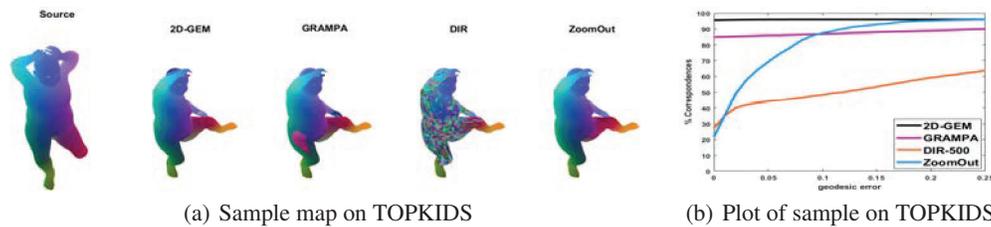


Figure 1: A non-isometric example. Illustrating the performance of 2D-GEM and other methods on a sample shape from the TOPKIDS dataset. For DIR we used DIR-500 with all LMD thresholds set to 5 since geodesic distances are not conserved on TOPKIDS (note that here, DIR performs much worse with lower thresholds). For ZoomOut we used 200 eigen-values and started from 20 with a step size of 1. For 2D-GEM and GRAMPA we followed the parameters in Section D.4

First, on non-isometric shapes, given that most existing methods assume near-isometry, they cannot handle noisy shapes (e.g., shapes with strong topological noise such as mesh “gluing” in areas of contact as in the TOPKIDS dataset [42]), since in this case geometric properties such as geodesic distances are not preserved and hence the shapes are not nearly isometric. An illustration can be seen in Figures 1(a) and 1(b), where the existing methods that assume near-isometry such as ZoomOut [54] and DIR [80] do not perform well at geodesic error 0 on the sample shape from TOPKIDS. They even produce a non-smooth map in the case of DIR. Moreover, those methods that try to address this noise are computationally very expensive and do not generalize well to nearly-isometric shapes, because for nearly-isometric shapes stronger constraints (such as preserving geodesic distances) are needed for the match. To illustrate this point, the spectral method GRAMPA is very expensive as it needs to compute all the eigen-values and eigen-vectors of the adjacency matrices obtained from the triangulations of each shape (see more details in Appendix B and time comparison with 2D-GEM in Figures 2(a) and 2(b)). Moreover, GRAMPA cannot generalize to cases of near-isometry

as illustrated in Figure 3(a), which shows the matching of GRAMPA in an example from SCAPE, a nearly-isometric shape dataset).

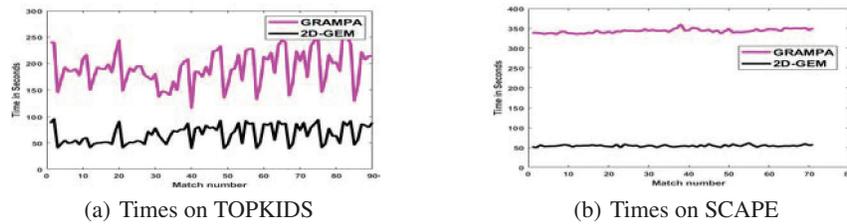


Figure 2: Time taken by 2D-GEM and GRAMPA for each of the 90 shapes on TOPKIDS and 71 on SCAPE.

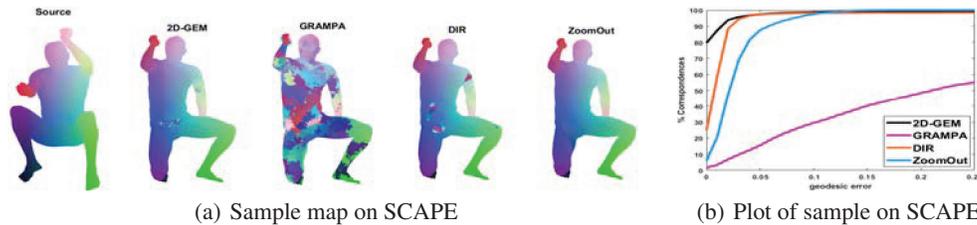


Figure 3: An isometric example. Illustrating the performance of 2D-GEM and other methods on a sample shape from the SCAPE dataset (using same settings as in figure 1).

Second, for nearly-isometric shapes, most existing methods attempt to first find an initial correspondence and then iteratively improve this initial correspondence using all points and the eigen-vectors corresponding to lower frequencies (eigen-values) of the LBO [4, 58, 54] or using other shape descriptors [64]. However, such an approach leads to an accumulation of errors in the correspondence, since the initial correspondence is mostly not accurate. This problem has been discussed in details by Xiang et al. [80], which attempted to address this via the Dual Iterative Refinement (DIR). DIR finds well matched points using the notion of Local Mapping Distortion (LMD) to identify points for which the geodesic distances are well preserved by the correspondence, and then use these points to update the correspondences of the other points via the notion of Functional maps [56]. Thus, DIR reduces the error accumulation. However, the DIR approach only works for nearly-isometric shapes and not for non-isometric shapes since it is based on the assumption of preserving geodesic distances which are not preserved in the case of non-isometries (see Figure 1(a), which shows that DIR does not perform well at geodesic 0 for non-isometric shapes). Finally, even on nearly-isometric shapes, DIR, though effective, requires a lot of eigen-vectors of the LBO in order to perform well at geodesic error 0. This can be seen in Figure 3(b), where using up to 500 eigen-vectors gives only 25% accuracy at geodesic error 0 on the sample shape from SCAPE, and even with up to 1000 eigen-vectors it only achieves 69.5% correspondence accuracy on SCAPE and a 59.8% on TOSCA at geodesic error 0 (as reported in the experiments in Section 5).

Given these problems, we propose to utilize 2D graph convolution for (a) non-isometric shape matching (since it is a de-noising operation), and (b) for isometric shape matching by further constraining our framework to preserve geodesic distances.

To this end:

- we show how graph matching can be done via the 2D graph convolution, and re-define noise in the context of graph matching. We further show empirically that this approach can match the recent spectral method GRAMPA on Erdos Regny random graph matching, as well as outperform it in the presence of informative node features (see Appendix D),
- we propose a new 2D graph convolution-based approach to shape matching, called 2D-GEM, which is (a) robust to noise, and achieves high accuracy on exact matching of non-isometric shapes while being orders of magnitude faster than GRAMPA, and (b) with some additional constraints also achieves high accuracy on exact matching of nearly-isometric shapes.

2 2D graph Convolution for Joint Alignment

To the best of our knowledge, we are the first to propose a 2-D graph convolution for graph alignment and shape matching. The most closely related work is the work that uses graph convolution neural networks for graph alignment [22] (See Appendix A.2) Here, we revisit the traditional graph convolution, show how the *joint alignment* problem can be done via 2D graph convolution, and finally we re-define noise in the context of graph matching.

2.1 Graph Convolution

From the Graph Signal Processing [67] perspective, graph convolution is regarded as a frequency filtering operator over node attributes. Given a graph, with node attribute matrix $\mathbf{X} \in \mathbb{R}^{n \times d}$, its normalized symmetric Laplacian (referred to as Laplacian in the rest of the paper) can be eigen-decomposed as $\mathbf{L} = \mathbf{U}\Lambda\mathbf{U}^T$, where $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_n) \in \mathbb{R}^{n \times n}$ are the eigen-values in increasing order, and $\mathbf{U} = [\mathbf{u}_1, \dots, \mathbf{u}_n] \in \mathbb{R}^{n \times n}$ are the associated orthogonal eigen-vectors. The graph convolution operation applied to the graph signal \mathbf{x} (a column of \mathbf{X}) consists of three steps [45, 24, 29, 30]: (1) transform node attributes into the spectral domain (i.e., $\mathbf{z} = \mathbf{U}^T\mathbf{x}$); (2) remove attribute noise by scaling the spectrum with a function defined on eigen-values (i.e., $\bar{z}_i = g(\lambda_i)z_i$); (3) perform inverse transformation on the scaled spectrum (i.e., $\bar{\mathbf{x}} = \mathbf{U}\bar{\mathbf{z}}$), where $g(\lambda)$ is a scalar-valued function on the eigen-values λ . In short, graph convolution is able to remove the noisy signals by assigning smaller weight on higher eigen-value spectrum with a well defined function $g(\lambda)$. For convenience, we denote $g(\mathbf{L}) = \mathbf{U}g(\Lambda)\mathbf{U}^T$, so that we have the convolution as $g * \mathbf{X} = g(\mathbf{L})\mathbf{X}$. We refer to $g(\mathbf{L})$ as the *convolution function* and $g(\Lambda)$ as the *filter function* or simply filter.

2.2 Joint Alignment via 2D graph Convolution

Following previous works [22, 71] and Appendix C, the graph convolution is invariant to permutation. As such, given two graphs \mathcal{A} and \mathcal{B} with graph convolution functions $\mathbb{M}_A = g(\mathbf{L}_A), \mathbb{M}_B = h(\mathbf{L}_B) \in \mathbb{R}^{n \times n}$, features $\mathbf{X}_A, \mathbf{X}_B \in \mathbb{R}^{n \times d}$, and having the functions g and h to be the same function, the *joint alignment* problem in Equation (1) can instead be performed as:

$$\min_{\mathbf{P}} \|\mathbf{P}\mathbb{M}_A\mathbf{X}_A - \mathbb{M}_B\mathbf{X}_B\|_F^2, \quad (2)$$

see Appendix C for more details. We observe that (2) reduces to the Linear Assignment Problem (LAP):

$$\max_{\mathbf{P}} \text{Tr}(\mathbf{P}\mathbb{M}_A\mathbf{X}_A\mathbf{X}_B^T\mathbb{M}_B), \quad (3)$$

where $\text{Tr}(\cdot)$ is the matrix trace operator, the convolution functions capture the pair-wise descriptors, and the features capture the point-wise descriptors. The first benefit of the convolution is to help do the joint alignment as a Linear Assignment Problem (LAP), which is much cheaper than the traditional Quadratic Assignment Problem (QAP) [23] used for graph matching. Note that, $\mathbb{M}_A = \mathbf{U}g(\Lambda)\mathbf{U}^T = \sum_i g(\lambda_i)\mathbf{u}_i\mathbf{u}_i^T$ and $\mathbb{M}_B = \mathbf{V}h(\mathbf{M})\mathbf{V}^T = \sum_j h(\mu_j)\mathbf{v}_j\mathbf{v}_j^T$, where \mathbf{U}, \mathbf{V} are the orthogonal eigen-vectors of \mathbb{M}_A and \mathbb{M}_B , and Λ, \mathbf{M} are the diagonal matrices of their respective eigen-values. Given these, one can rewrite our cost matrix $\mathbf{C} = \mathbb{M}_A\mathbf{X}_A\mathbf{X}_B^T\mathbb{M}_B$ in (3) as:

$$\mathbf{C} = \sum_{i,j} g(\lambda_i)h(\mu_j)\mathbf{u}_i\mathbf{u}_i^T\mathbf{X}_A\mathbf{X}_B^T\mathbf{v}_j\mathbf{v}_j^T, \quad (4)$$

When node features are not available, one can initialize all nodes with the same features in (3), i.e., $\mathbf{X}_A = \mathbf{J}_A \in \mathbb{R}^{n_1 \times d}$, and $\mathbf{X}_B = \mathbf{J}_B \in \mathbb{R}^{n_2 \times d}$, where n_1 may or may not be equal to n_2 for different size graphs, and $\mathbf{J}_A, \mathbf{J}_B$ are all one rectangular matrices. This is equivalent to assuming that nodes have the same probability of being matched to each other. Moreover, to save cost in the case when there is no feature, one could simply initialize the similarity matrix $\mathbf{X}_A\mathbf{X}_B^T = \mathbf{J}_{A,B} \in \mathbb{R}^{n_1 \times n_2}$ (given that $\mathbf{J}_A\mathbf{J}_B^T = \mathbf{J}_{A,B} \in \mathbb{R}^{n_1 \times n_2}$, where $\mathbf{J}_{A,B}$ is a scalar matrix with constant $[\mathbf{J}_{A,B}]_{ij} = \sum_{i=1}^d 1 = d$) for all entries. In the case where there is no feature, (4) becomes:

$$\mathbf{C} = \sum_{i,j} g(\lambda_i)h(\mu_j)\mathbf{u}_i\mathbf{u}_i^T\mathbf{J}_{A,B}\mathbf{v}_j\mathbf{v}_j^T. \quad (5)$$

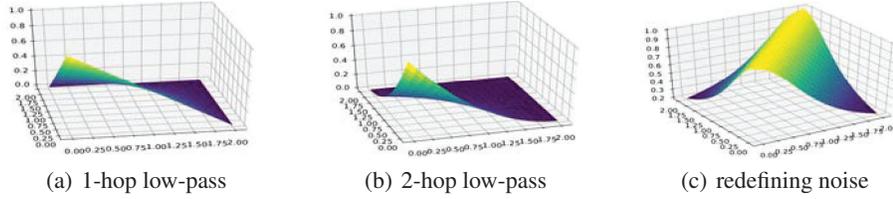


Figure 4: 2D graph filters. Figure 4(a) uses the convolution functions $g(\mathbf{L}_A) = g(\mathbf{L}_B) = (\mathbf{I} - \frac{1}{2}\mathbf{L})$, Figure 4(b) uses the convolution functions $g(\mathbf{L}_A) = g(\mathbf{L}_B) = (\mathbf{I} - \frac{1}{2}\mathbf{L})^2$, and Figure 4(c) uses the noise robust filter defined in Section 3 with $e = 1$.

Our cost matrices in equations (4) and (5) can be seen as 2D graph convolution on the 2D graph signals $\mathbf{X}_A \mathbf{X}_B^T$ (or $\mathbf{J}_{A,B}$ in the case of no features) [46], where we: (1) transform the 2D graph signals into the spectral domain to obtain the 2D graph spectrum (i.e., $\mathbf{Z} = \mathbf{U}^T \mathbf{X}_A \mathbf{X}_B^T \mathbf{V}$ or $\mathbf{Z} = \mathbf{U}^T \mathbf{J}_{A,B}^T \mathbf{V}$) [46]; (2) remove noise (de-noise) from the 2D graph signals by scaling the 2D spectrum with a function defined on each pair of eigen-values (i.e., $\tilde{z}_{ij} = w(\lambda_i, \mu_j) z_{ij}$); and (3) perform inverse transformation on the scaled 2D spectrum ($\mathbf{C} = \mathbf{U} \tilde{\mathbf{Z}} \mathbf{V}^T$). Thus, our cost matrices are 2D filtering operation on the 2D graph signals, where $w(\lambda_i, \mu_j) = g(\lambda_i)h(\mu_j)$ is the 2D filter (see Figures 4(a), 4(b) and 4(c) for different 2D filters).

3 Redefining Noise in the Context of Graph Matching

In this section, we redefine noise in the context of graph matching and propose a suitable filter for graph convolution with regards to the graph matching task. It can be shown that graph convolution is a solution to the graph signal de-noising problem [24]. Formally, assuming that (1) the ground-truth node features $\hat{\mathbf{X}}$ are smooth w.r.t. a graph \mathcal{G} with adjacency \mathbf{A} and Laplacian \mathbf{L} , and that (2) the node feature noise has small magnitude, then the total variation of $\hat{\mathbf{X}}$ should be small by assumption (1), and $\|\hat{\mathbf{X}} - \mathbf{X}\|$ can be upper-bounded by assumption (2). Thus, we can model the problem of recovering the ground-truth signal as:

$$\hat{\mathbf{X}}^* = \arg \min_{\hat{\mathbf{X}}} \text{Tr} \left(\hat{\mathbf{X}}^T \mathbf{L} \hat{\mathbf{X}} \right) \quad \text{s.t.} \|\hat{\mathbf{X}} - \mathbf{X}\|_F^2 \leq \epsilon_1, \quad (6)$$

It can clearly be observed that $\hat{\mathbf{X}}^* = (\mathbf{I} + \alpha \mathbf{L})^{-1} \mathbf{X}$, where α is a hyperparameter. This is a graph convolution with graph convolution function given by $(\mathbf{I} + \alpha \mathbf{L})^{-1}$. This function can be decomposed as $\mathbf{U}(\mathbf{I} + \alpha \Lambda)^{-1} \mathbf{U}^T$. As well studied in [45, 78, 36], one can see that the filter $(\mathbf{I} + \alpha \Lambda)^{-1}$ is decaying, i.e., $g(\lambda_i)$ is lower for higher λ_i , which will mostly reduce the amplitude z_i of graph signals \mathbf{u}_i belonging to larger eigen-values λ_i . Thus, one can conclude that noise corresponds to higher order frequencies (eigen-values). This kind of filter is commonly known as a low-pass filter [45]. It can be shown that for the 2D graph convolution, filters of the form $w(\lambda_i, \mu_j) = g(\lambda_i)h(\mu_j)$, where $g(\lambda_i) = (1 - \alpha \lambda_i)^k$ and $h(\mu_j) = (1 - \alpha \mu_j)^k$ will result in a 2D low-pass filter as well (see Figures 4(a) and 4(b)). These filters are referred to as k -hop low-pass filters (aggregating information from the k -hop neighborhood of a node).

However, in the context of graph matching, graphs are isomorphic when they have the same eigen-value, and in this case we want to align the eigen-vectors since two graphs \mathcal{A} and \mathcal{B} are isomorphic iff their eigen-vectors are permutations of each other up to a sign [72, 37].

This leads us to redefine noise in the graph matching context. Given two graphs \mathcal{A} and \mathcal{B} , with convolution functions $\mathbb{M}_A, \mathbb{M}_B$ whose eigen-values are λ_i and μ_j respectively, noise in graph matching corresponds to 2D graph signals $z_{i,j} = \mathbf{u}_i^T \mathbf{X}_A \mathbf{X}_B^T \mathbf{v}_j$ or $z_{i,j} = \mathbf{u}_i^T \mathbf{J}_{A,B} \mathbf{v}_j$ such that $\lambda_i \neq \mu_j$, i.e., eigen-values that differ. Thus, we should use a filter that kills off these signals corresponding to eigen-values that do not match. We choose the filter $w(\lambda_i, \mu_j, e) = \frac{1}{(\lambda_i^e - \mu_j^e)^2 + 1}$ instead of $g(\lambda_i)h(\mu_j)$. This filter exactly kills off these signals corresponding to eigen-values that differ (see Figure 4(c) to see an illustration of its action on signals). A note here is that the amplitude and the width of the filter response in Figure 4(c) depends on the spectra of the graphs to be matched as well as on e . For simplicity, we denote this filter as $w(\lambda_i, \mu_j)$ in the rest of the paper. Thus, a

Algorithm 1 : 2D-GEM

1. **input** An initial correspondence \mathbf{P}^0 , \mathbf{U}_k , \mathbf{V}_k , $\mathbf{\Lambda}_k$, \mathbf{M}_k , e , maximum iteration $iter$, and the LMD error $\{\epsilon_i\}_{i=1}^{iter}$ thresholds.
while $0 \leq i \leq iter$ **do**
 2. use one round of the 2-Hop algorithm. i.e., set $t = 1$ in Algorithm 2
 3. use the LMD to locate well anchored pairs (landmarks) l
 4. use these well anchored pairs to update the 2D graph spectrum as in (9)
 5. use the updated spectrum for updating the correspondence of poorly matched points p :
 $\max_{\mathbf{P}_p} \text{Tr}(\mathbf{P}_p \mathbf{C}_p)$, with \mathbf{C}_p as in (10)
end while

Algorithm 2 : 2-Hop

1. **input** An initial correspondence \mathbf{P} , 2-hop adjacency matrices \mathbf{A}_2 , \mathbf{B}_2 , maximum iteration t .
while $0 \leq i \leq t$ **do**
 2. $\mathbf{P}^{i+1} = GMWM(\mathbf{W})$, where $\mathbf{W} = \mathbf{A}_2 \mathbf{P}^i \mathbf{B}_2$
end while
3. **return** \mathbf{P}^t

noise robust 2D graph convolution based cost matrix for graph matching is given as:

$$\mathbf{C} = \sum_{i,j} w(\lambda_i, \mu_j) \mathbf{u}_i \mathbf{u}_i^T \mathbf{X}_A \mathbf{X}_B^T \mathbf{v}_j \mathbf{v}_j^T. \quad (7)$$

Moreover, to reduce the time complexity, one may use the first k eigen-values $\mathbf{\Lambda}_k = \text{diag}(\lambda_1, \dots, \lambda_k)$, $\mathbf{M}_k = \text{diag}(\mu_1, \dots, \mu_k)$ and eigen-vectors $\mathbf{U}_k = [\mathbf{u}_1, \dots, \mathbf{u}_k]$, $\mathbf{V}_k = [\mathbf{v}_1, \dots, \mathbf{v}_k]$, instead of the full spectrum. Doing so, the noise robust 2D graph convolution based cost matrix for graph matching becomes:

$$\mathbf{C} = \sum_{i,j}^k w(\lambda_i, \mu_j) \mathbf{u}_i \mathbf{u}_i^T \mathbf{X}_A \mathbf{X}_B^T \mathbf{v}_j \mathbf{v}_j^T. \quad (8)$$

Due to space constraints, we refer the reader to Appendix D for the experimental validation of the 2D graph convolution approach for graph matching.

4 Our Algorithm

Having formulated the 2D graph convolution-based approach for graph matching as well as redefined noise in the context of graph matching, we now formulate 2D-GEM (2D Graph convolution-based framework for attempting Exact shape Matching) for non-isometric and nearly-isometric 3D shape matching.

4.1 2D-GEM

For our framework, given the xyz coordinates of two shapes \mathcal{M}_1 and \mathcal{M}_2 , we build suitable point-wise descriptors \mathbf{X}_A and \mathbf{X}_B , as well as graph laplacians \mathbf{L}_A and \mathbf{L}_B using their triangulations. We then initialize the correspondence using the point-wise descriptors. Then at each iteration, we first follow GRAMPA and update the correspondence using the 2-hop neighborhood graph matching algorithm [84], which makes use of the Greedy Maximum Weight Matching (GMWM) with computational complexity $O(n^2 \log n)$ (see Algorithm 2 and Appendix B). The 2-hop adjacency matrices \mathbf{A}_2 and \mathbf{B}_2 of the shapes \mathcal{M}_1 and \mathcal{M}_2 , respectively, are such that $[\mathbf{A}_2]_{i,j} = 1$ and $[\mathbf{B}_2]_{i,j} = 1$ if i and j are 2-hop neighbors, and 0 otherwise. Second, we use the notion of conservation of geodesic distances [64] to further constrain 2D-GEM, we capture this via LMD[80]. Finally, we adopt the 2D graph convolution in Equation (8) to filter the correspondence \mathbf{P} rather than the feature similarity matrix $\mathbf{X}_A \mathbf{X}_B^T$ in Equation (8), i.e., we use \mathbf{P} rather than $\mathbf{X}_A \mathbf{X}_B^T$ as our 2D graph signals. This can be seen as de-noising the correspondence at each iteration following the state-of-the-art Kernel-Matching method [42].

In details (see Algorithm 1), we initialize the correspondence via SHOT[70] (though other descriptors and algorithms could also be used as in Appendices D, and F), and then at each iteration ($iter$):

- we update the correspondence using the 2-Hop Algorithm (2) with $t = 1$,
- we find l anchor pairs and p non-anchor pairs via LMD (see Section 4.2) and use them to update the 2D graph spectrum as

$$\mathbf{Z}_l = \mathbf{U}_k^T[l] \mathbf{P}_l \mathbf{V}_k[l], \quad (9)$$

where $\mathbf{U}_k[l], \mathbf{V}_k[l] \in \mathbb{R}^{l \times k}$ and $\mathbf{P}_l \in \mathbb{R}^{l \times l}$ are the eigen-vectors and correspondence of the l -anchor pairs (landmarks) found, and \mathbf{Z}_l is the updated $k \times k$ 2D graph spectrum.

- we use the updated spectrum for updating the correspondence of poorly matched points as $\max_{\mathbf{P}_p} \text{Tr}(\mathbf{P}_p \mathbf{C}_p)$, where

$$\mathbf{C}_p = \sum_{i,j}^k w(\lambda_i, \mu_j) \mathbf{u}_i[p] [\mathbf{Z}_l]_{i,j} \mathbf{v}_j^T[p], \quad (10)$$

and $\mathbf{u}_i[p], \mathbf{v}_j[p] \in \mathbb{R}^p$ are the eigen-vectors of the non-landmark points p .

Computation complexity. The complexity for computing k leading eigen-vectors of a $n \times n$ sparse matrix is $O(kn \log n)$. The complexity for checking LMD is $O(n)$ following DIR [80]. The GMWM has a computational complexity of $O(n^2 \log n)$. Thus, our method is of complexity $O(\text{iter} n^2 \log n) = O(n^2 \log n)$, since *iter* is usually small, e.g., 10 or 60 in our experiments.

4.2 Local Mapping Distortion (LMD)

Here, we describe the LMD employed by 2D-GEM to get landmarks. The LMD was introduced by [80] and [79] for nearly-isometric shapes. Let $T : \mathcal{M}_1 \rightarrow \mathcal{M}_2$ be a map function between two shapes. The LMD of the map T at the point x_i is given as follows:

$$\mathcal{D}_\gamma(T)(x_i) = \frac{\sum_{x_j \in \mathcal{B}_\gamma(x_i)} \mathcal{A}_1(j) DE_T(x_i, x_j)}{\sum_{x_j \in \mathcal{B}_\gamma(x_i)} \mathcal{A}_1(j)}, \quad (11)$$

where $\mathcal{B}_\gamma(x) = \{y \in \mathcal{M}_1 \mid d_{\mathcal{M}_1}(x, y) \leq \gamma\}$ is the γ -geodesic ball of x , \mathcal{A}_1 is the area element of the mesh of shape \mathcal{M}_1 , and $DE_T(x, y) = |d_{\mathcal{M}_1}(x, y) - d_{\mathcal{M}_2}(T(x), T(y))|/\gamma$ represents a pair-wise distance distortion of mapping nearby points x_i and x_j to $T(x_i)$ and $T(x_j)$. A smaller value of $\mathcal{D}_\gamma(T)(x_i)$ means a better map continuity of T at the point x_i , in other words, the local distance at the point x_i is well preserved. Based on the above definition of LMD, one can check that if T is an isometric map, then $\mathcal{D}_\gamma(T)(x) = 0, \forall x \in \mathcal{M}_1, \gamma > 0$. Conversely, if $\mathcal{D}_\gamma(T)(x) = 0, \forall x \in \mathcal{M}_1$ for some $\gamma > 0$, then T is isometric. We use the LMD to find well matched pairs, i.e., pairs for which $\{(x_i, T(x_i)) \mid \mathcal{D}_\gamma(T)(x_i) \leq \epsilon_i\}$, where ϵ_i is a threshold. We call them anchor points (landmarks l), and others non-anchor points (poor-points p). We use the landmarks to update the 2D graph spectrum \mathbf{S}_{kl} , which we use to update the correspondence of non-landmark points (see Section 4.1). It is important to note that at every iteration we gradually decrease the threshold following Xiang et al. [80], and hence not only is the correspondence T updated, but also the landmarks l . For non-isometric shapes we relax this constraint by setting the thresholds ϵ_i arbitrarily large, such as 100 for 2D-GEM in the experiments (see Section 5.5).

4.3 Relationship Between 2D-GEM and Functional Maps

Here we show the relationship between 2D-GEM and the popular Functional Map [56]. The notion of Functional Maps [56] was introduced in order to solve the notoriously difficult QAP resulting from *pair-wise alignment*. A Functional Map is a map between functional spaces $P : L^2(\mathcal{M}_1) \rightarrow L^2(\mathcal{M}_2)$, which can be discretized as an $n \times n$ matrix \mathbf{P} (assuming that each of the shapes \mathcal{M}_1 and \mathcal{M}_2 has n vertices). Given a pair of orthonormal basis $\mathbf{U} = [\mathbf{u}_1, \dots, \mathbf{u}_n]$ and $\mathbf{V} = [\mathbf{v}_1, \dots, \mathbf{v}_n]$ for $L^2(\mathcal{M}_1)$ and $L^2(\mathcal{M}_2)$, respectively, we have $\mathbf{P} = \mathbf{U} \mathbf{S} \mathbf{V}^T$, where \mathbf{S} acts as a basis transformation matrix. On the other hand, recalling 2D-GEM's cost matrix for iteratively de-noising the Permutation matrix in Algorithm 1, i.e.,

$$\mathbf{C} = \sum_{i,j}^k w(\lambda_i, \mu_j) \mathbf{u}_i \mathbf{u}_i^T \mathbf{P} \mathbf{v}_j \mathbf{v}_j^T = \sum_{i,j}^k w(\lambda_i, \mu_j) \mathbf{u}_i \mathbf{S}_{i,j} \mathbf{v}_j^T, \quad (12)$$

where $s_{i,j} = \mathbf{u}_i^T \mathbf{P} \mathbf{v}_j$. This means that one can interpret the matrix $\mathbf{S}_k \in \mathbb{R}^{k \times k}$ containing the $s_{i,j}$'s as a low-pass approximation of the permutation matrix \mathbf{P} in the truncated Laplacian eigen-basis, i.e., $\mathbf{P} \approx \mathbf{U} \mathbf{S}_k \mathbf{V}^T$ [42]. Equation (12) can thus be interpreted as applying a graph matching-suitable filter to the functional map matrix \mathbf{S}_k [56, 42]. 2D-GEM 1 can thus be regarded as denoising the functional map in the graph laplacian basis while enforcing the preservation of neighborhood using 2-hop witnesses [84], as well as ensuring the preservation of geodesic distances (distortions) via the LMD [80]. This interpretation leads to the beautiful connection between the 2D-graph spectrum and the functional map. Essentially they are the same when we take the 2D-graph signals as the columns of the correspondence matrix. As such the 2D-graph convolution approach is intricately connected to methods such as ZoomOut [54], ICP [4] and BCICP [58], which use the functional map together with the truncated Laplacian eigen-basis to update the correspondence \mathbf{P} . Moreover, recent works such as Structured Regularization of Functional Map Computations [59], and DIR[80] tried to address some problems with the traditional functional map approaches (namely, (a) penalizing the functional map update only when it fails to commute with the LBO's of the source and target shapes often fails to produce a good functional map, (b) accumulation of errors in the map update due to poor initial correspondence) by introducing regularizers in the functional map update step. Following these works, 2D-GEM can also be seen as a framework for further constraining the functional map and the point-wise map updates using the LMD and the 2-hop neighborhood witnesses. Another line of work, namely Fast Sinkhorn Filters [57] tried to explain functional map based approaches as spectral alignment regularized algorithms, and as such, 2D-GEM can also be viewed as a spectral alignment regularized algorithm.

5 Experiments

We report experimental results that validate the effectiveness and efficiency of 2D-GEM on exact matching of non-isometric and nearly-isometric 3D shapes.

5.1 Experimental Set-up

Time comparison between 2D-GEM and GRAMPA is conducted in Matlab on a Windows 10 system with 16GB RAM and Intel(R) i5 11400 CPU @ 2.60-4.4GHz.

5.2 Datasets

We conduct experiments to evaluate the performance of our method on two nearly isometric benchmark datasets TOSCA [10], and SCAPE [5], as well as a benchmark non-isometric dataset SHREC'16 (TOPKIDS) [42]. TOSCA consists of 80 shapes in 8 different categories (human and animal shapes) with vertex numbers ranging from 4k to 50k. SCAPE has 71 shapes (12,500 vertices for each) of the same person with different poses. TOPKIDS contains 25 shapes of the same class with up to 12K vertices, undergoing near-isometric deformations in addition to large topological shortcuts (noise).

5.3 Evaluation Metrics

We use the geodesic error as our error metric [20]. Given that the correspondence of an algorithm maps $x \in \mathcal{M}_1$ to $y \in \mathcal{M}_2$, and the true correspondence maps x to y^* , the geodesic error is defined as $e(x) = \frac{d_{\mathcal{M}_2}(y, y^*)}{diam(\mathcal{M}_2)}$, where $d_{\mathcal{M}_2}$ denotes the geodesic distance on \mathcal{M}_2 , and $diam(\mathcal{M}_2)$ is the geodesic diameter of \mathcal{M}_2 .

5.4 Baselines

For non-isometric shape matching, i.e., for matching on the TOPKIDS dataset, we compare 2D-GEM with the following methods: EM [62], GE [41], RF [61], PFM [60], FSPM [49], Kernel-Matching [42], and GRAMPA [20].

For near-isometric shape matching, i.e., for matching on the SCAPE and TOSCA datasets, we compare 2D-GEM to the following methods: SGMDS[2], FM[56], BIM[34], Mobius[48], Best-Conformal [34], Kernel-Matching [42], DIR-500 [80] which uses 500 eigen-vectors, DIR-1000 [80]

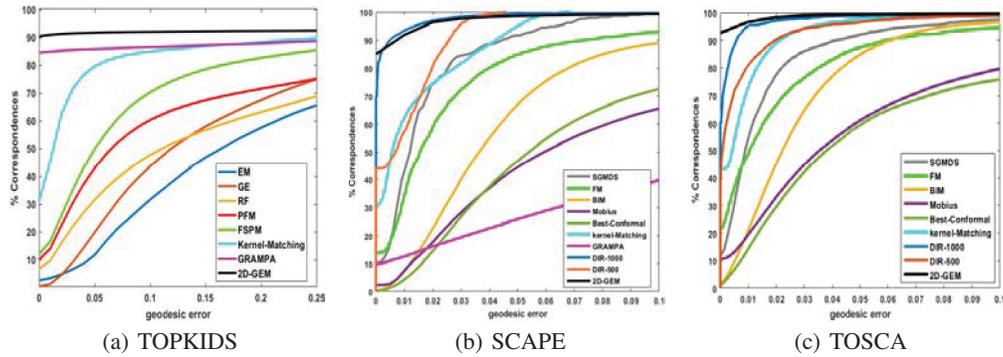


Figure 5: Performance comparison on TOPKIDS 5(a), SCAPE 5(b), and TOSCA 5(c).

which uses 1000 eigen-vectors, and GRAMPA (on SCAPE, not on TOSCA due to its time complexity and the fact that 2D-GEM significantly outperforms it as shown on TOPKIDS and SCAPE).

5.5 Parameter Settings

For 2D-GEM, we initialize the correspondence by the GMWM on SHOT [70], and use $e = 10$ (see Appendix D for more analysis on e). For nearly-isometric shapes (TOSCA and SCAPE), we use 500 eigen-vectors of the LBO. Following Xiang et al. [80], we use the second ring neighborhood size for LMD criterion and set the maximum iteration number to $iter = 10$. We also set the LMD thresholds, we use $[0.5280, 0.4560, 0.3840, 0.3120, 0.2400, 0.1680, 0.1680, 0.1680, 0.1680, 0.1680]$. For non-isometric shapes (TOPKIDS), we set the thresholds to 100 for all iterations, use $iter = 60$, and 10 LBO eigen-vectors.

For other baselines, we report their results as reported in their papers i.e., we use their curves. Except for GRAMPA where we re-run using the 2-hop improvement algorithm instead which performs much better than the 1-hop, i.e., the adjacency. For fair comparison with 2D-GEM, we set $t = 60$ for this improvement on GRAMPA.

5.6 Performance Analysis

Comparison on non-isometric shapes.² GRAMPA is the state-of-the-art non-isometric shape matching method that achieves an excellent performance on the exact matching task on the noisy dataset TOPKIDS (Figure 5(a)). It achieves 84.6% on this dataset at geodesic error 0. This shows that GRAMPA is very robust to noise as expected (note that this performance is much higher than the 60% reported in their original paper because we used the 2-hop adjacency instead of the 1-hop for refinement as explained in Appendix B.2). However, our method 2D-GEM further outperforms GRAMPA and all other methods, including GRAMPA at geodesic error 0. As reported in Figure 5(a), where 2D-GEM achieves a 90.5% accuracy. Moreover, Figures 2(a) and 2(b) show that 2D-GEM is significantly faster than GRAMPA since 2D-GEM only needs to compute a subset of the eigen-vectors, unlike GRAMPA which needs the whole spectrum and thus has a complexity of $O(n^3)$.

Comparison on nearly-isometric shapes. DIR is the state-of-the-art model for nearly-isometric shape matching. DIR outperforms all the other baselines except 2D-GEM. The better DIR-1000 using 1000 eigen-vectors achieves 69.5% accuracy at geodesic error 0 on SCAPE, and 59.8% accuracy at geodesic error 0 on TOSCA. Using the LMD proposed for the DIR model by Xiang et al. [80] for extra constraints on 2D-GEM, 2D-GEM not only generalizes to nearly-isometric shapes, but it achieves the state-of-the-art results using only 500 eigen-vectors, outperforming even DIR-1000. 2D-GEM achieves 92.5% accuracy (outperforming DIR-1000 by more than 40%) at geodesic error 0 on the TOSCA dataset, as well as 84.9% accuracy on the SCAPE dataset (outperforming DIR-1000 by almost 20%). Note that these 2D-GEM’s accuracy improves using more eigen-vectors or increasing $iter$, but this also increases the running time. Moreover, 2D-GEM’s gain over GRAMPA (the state-of-the-art non-isometric shape matching method) can also be seen to extend to nearly-isometric shapes. As noted earlier in Section 1, GRAMPA does not generalize to nearly-isometric shapes since the

²Ablation studies on 2D-GEM in Appendix E. Code at: <https://github.com/BarakeelFanseu/2D-GEM>

assumptions for nearly-isometric shape matching are different from those of non-isometric shape matching given that isometries preserve geometric properties which are lost by non-isometries. This is seen by its poor performance on the SCAPE dataset (Figure 5(b), where even for exact matching, i.e., at geodesic error 0, it only achieves around 10%). In contrast, 2D-GEM (using the LMD for extra constraints) generalizes to nearly-isometric shapes and achieves the state-of-the-art performance. We do not attempt to use the LMD on GRAMPA since (1) it is not clear how to use GRAMPA in an iterative way as it is usually only used as an initialization for other algorithms (see Appendix B), (2) 2D-GEM outperforms GRAMPA even without the LMD as shown on the noisy TOPKIDS dataset (Figure 5(a)), and (3) 2D-GEM with or without the LMD is still faster than GRAMPA without the LMD (see Figures 2(a) and 2(b)).

5.7 Limitations of 2D-GEM

There are three main limitations to this work. First, for nearly-isometric shapes (shown in Figure 3(a)), 2D-GEM sometimes produces maps that are not continuous (smooth), though much more accurate than the other baselines. Second, although the time complexity $O(n^2 \log n)$ of 2D-GEM is much lower than the $O(n^3)$ of GRAMPA, it is much worse than the $O(n \log n)$ of other methods such as ZoomOut and DIR. Lastly, unlike GRAMPA, we do not provide rigorous theoretical exact recovery guarantees for 2D-GEM on random graph matching.

6 Conclusions

We presented how the 2D graph convolution could be used for graph matching. Using this 2D graph convolution, we re-defined noise in the context of graph matching and proposed a model (2D-GEM) that addresses the task of exact 3D shape matching on non-isometric (noisy) datasets as well as on nearly-isometric datasets. We then showed that 2D-GEM significantly outperforms existing methods on both non-isometric shapes and nearly-isometric shapes. For future work, we propose to address two key limitations of 2D-GEM, i.e., further reducing the time complexity of $O(n^2 \log n)$ and providing theoretical guarantees on its exact recovery performance. Moreover, we also propose to explore the generalization of 2D-GEM for correspondence tasks that allow using graph representation such as partial shape matching, image matching, and network de-anonymization. Besides, it is also promising to consider the potential distribution shifts [12], as well as the optimization trade-offs [13].

7 Broader Impact

Given that our method needs no training, is unsupervised and based on intrinsic shape properties, it can positively benefit the industry and society as a whole in that; (a) it may be easily generalized, and most importantly (b) it needs no human labels nor training, which can help reduce the cost and time taken for the shape correspondence problems in different crucial fields such as medicine, and computer graphics.

8 Acknowledgements

We thank the reviewers for their valuable comments. This work was partially supported by GRF 14208318 from the RGC of HKSAR and the CUHK direct grant 4055146. Lin Zhang and Bo Li were supported in part by RGC RIF grant R6021-20 and RGC GRF grant under contracts 16209120 and 16200221. Bo Han was supported by NSFC YSF No. 62006202, RGC ECS No. 22200720, and Guangdong Basic and Applied Basic Research Foundation No. 2022A1515011652.

References

- [1] Yonathan Aflalo, Alexander Bronstein, and Ron Kimmel. On convex relaxation of graph isomorphism. *Proceedings of the National Academy of Sciences*, 112(10):2942–2947, 2015. ISSN 0027-8424. doi: 10.1073/pnas.1401651112.
- [2] Yonathan Aflalo, Anastasia Dubrovina, and Ron Kimmel. Spectral generalized multi-dimensional scaling. *International Journal of Computer Vision*, 118:380–392, 2016.

- [3] H.A. Almohamad and S.O. Duffuaa. A linear programming approach for the weighted graph matching problem. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(5): 522–525, 1993. doi: 10.1109/34.211474.
- [4] B. Amberg, S. Romdhani, and T. Vetter. Optimal step nonrigid icp algorithms for surface registration. In *2007 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, June 2007.
- [5] Dragomir Anguelov, Praveen Srinivasan, Hoi-Cheung Pang, Daphne Koller, Sebastian Thrun, and James Davis. The correlated correspondence algorithm for unsupervised registration of nonrigid surfaces. In *Advances in neural information processing systems*, pages 33–40, 2005.
- [6] Mathieu Aubry, Ulrich Schlickewei, and Daniel Cremers. The wave kernel signature: A quantum mechanical approach to shape analysis. In *2011 IEEE International Conference on Computer Vision Workshops (ICCV Workshops)*, pages 1626–1633, 2011. doi: 10.1109/ICCVW.2011.6130444.
- [7] László Babai, D. Yu. Grigoryev, and David M. Mount. Isomorphism of graphs with bounded eigenvalue multiplicity. In *Proceedings of the Fourteenth Annual ACM Symposium on Theory of Computing*, STOC '82, page 310–324, New York, NY, USA, 1982. Association for Computing Machinery. ISBN 0897910702. doi: 10.1145/800070.802206.
- [8] Yunsheng Bai, Hao Ding, Yizhou Sun, and Wei Wang. Convolutional set matching for graph similarity. In *NeurIPS Workshop on Relational Representation Learning*, 2018.
- [9] Yunsheng Bai, Haoyang Ding, Song Bian, Ting Chen, Yizhou Sun, and Wei Wang. Simgnn: A neural network approach to fast graph similarity computation. *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*, 2019.
- [10] Alexander M Bronstein, Michael M Bronstein, and Ron Kimmel. *Numerical geometry of non-rigid shapes*. Springer Science & Business Media, 2008.
- [11] M. M. Bronstein and I. Kokkinos. Scale-invariant heat kernel signatures for non-rigid shape recognition. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1704–1711, 2010.
- [12] Yongqiang Chen, Yonggang Zhang, Yatao Bian, Han Yang, Kaili Ma, Binghui Xie, Tongliang Liu, Bo Han, and James Cheng. Invariance principle meets out-of-distribution generalization on graphs. *arXiv preprint arXiv:2202.05441*, 2022.
- [13] Yongqiang Chen, Kaiwen Zhou, Yatao Bian, Binghui Xie, Kaili Ma, Yonggang Zhang, Han Yang, Bo Han, and James Cheng. Pareto invariant risk minimization. *arXiv preprint arXiv:2206.07766*, 2022.
- [14] Donatello Conte, Pasquale Foggia, Carlo Sansone, and Mario Vento. Thirty years of graph matching in pattern recognition. *IJPRAI*, 18:265–298, 05 2004. doi: 10.1142/S0218001404003228.
- [15] Chris Ding, Tao Li, and Michael I. Jordan. Nonnegative matrix factorization for combinatorial optimization: Spectral clustering, graph matching, and clique finding. In *2008 Eighth IEEE International Conference on Data Mining*, pages 183–192, 2008. doi: 10.1109/ICDM.2008.130.
- [16] Jian Ding, Zongming Ma, Yihong Wu, and Jiaming Xu. Efficient random graph matching via degree profiles. *Probability Theory and Related Fields*, 179:29–115, 2020.
- [17] B. L. Douglas. The weisfeiler-lehman method and graph isomorphism testing. *arXiv: Combinatorics*, 2011.
- [18] Nadav Dym, Haggai Maron, and Yaron Lipman. Ds++: A flexible, scalable and provably tight relaxation for matching problems. *ACM Trans. Graph.*, 36(6), November 2017. ISSN 0730-0301. doi: 10.1145/3130800.3130826.
- [19] Frank Emmert-Streib, Matthias Dehmer, and Yongtang Shi. Fifty years of graph matching, network alignment and network comparison. *Information Sciences*, 346-347:180–197, 02 2016. doi: 10.1016/j.ins.2016.01.074.

- [20] Zhou Fan, Cheng Mao, Yihong Wu, and Jiaming Xu. Spectral graph matching and regularized quadratic relaxations: Algorithm and theory. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 2985–2995. PMLR, 13–18 Jul 2020.
- [21] Soheil Feizi, Gerald Quon, Mariana Recamonde-Mendoza, Muriel Medard, Manolis Kellis, and Ali Jadbabaie. Spectral alignment of graphs. *IEEE Transactions on Network Science and Engineering*, 7(3):1182–1197, jul 2020. ISSN 2327-4697. doi: 10.1109/TNSE.2019.2913233.
- [22] Matthias Fey, Jan E. Lenssen, Christopher Morris, Jonathan Masci, and Nils M. Kriege. Deep graph matching consensus. In *International Conference on Learning Representations*, 2020.
- [23] Gerd Finke, Rainer E. Burkard, and Franz Rendl. Quadratic assignment problems. In Silvano Martello, Gilbert Laporte, Michel Minoux, and Celso Ribeiro, editors, *Surveys in Combinatorial Optimization*, volume 132 of *North-Holland Mathematics Studies*, pages 61–82. North-Holland, 1987. doi: [https://doi.org/10.1016/S0304-0208\(08\)73232-8](https://doi.org/10.1016/S0304-0208(08)73232-8).
- [24] Guoji Fu, Yifan Hou, Jian Zhang, Kaili Ma, Barakeel Fansu Kamhoua, and James Cheng. Understanding graph neural networks from graph signal denoising perspectives. *ArXiv*, abs/2006.04386, 2020.
- [25] S. Gold and A. Rangarajan. A graduated assignment algorithm for graph matching. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18(4):377–388, 1996. doi: 10.1109/34.491619.
- [26] Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. In *Advances in neural information processing systems*, pages 1024–1034, 2017.
- [27] Mark Heimann, Haoming Shen, Tara Safavi, and Danai Koutra. Regal: Representation learning-based graph alignment. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, CIKM ’18, page 117–126, New York, NY, USA, 2018. Association for Computing Machinery. ISBN 9781450360142. doi: 10.1145/3269206.3271788.
- [28] Ruqi Huang, Jing Ren, Peter Wonka, and Maks Ovsjanikov. Consistent zoomout: Efficient spectral map synchronization. In *Proc. SGP*, volume 39, 2020.
- [29] Barakeel Fansu Kamhoua, Lin Zhang, Kaili Ma, James Cheng, Bo Li, and Bo Han. Hypergraph convolution based attributed hypergraph clustering. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, 2021.
- [30] Barakeel Fansu Kamhoua, Lin Zhang, Kaili Ma, James Cheng, Bo Li, and Bo Han. Grace: A general graph convolution framework for attributed graph clustering. *ACM Trans. Knowl. Discov. Data*, jun 2022. ISSN 1556-4681. doi: 10.1145/3544977. URL <https://doi.org/10.1145/3544977>. Just Accepted.
- [31] Ehsan Kazemi and Matthias Grossglauser. On the structure and efficient computation of isorank node similarities. *ArXiv*, abs/1602.00668, 2016.
- [32] Ehsan Kazemi, Seyed Hamed Hassani, M. Grossglauser, and H. Modarres. Proper: global protein interaction network alignment through percolation matching. *BMC Bioinformatics*, 17, 2016.
- [33] Itay Kezurer, Shahar Z. Kovalsky, Ronen Basri, and Yaron Lipman. Tight relaxation of quadratic matching. *Computer Graphics Forum*, 34(5):115–128, 2015. doi: <https://doi.org/10.1111/cgf.12701>.
- [34] Vladimir G. Kim, Yaron Lipman, and Thomas Funkhouser. Blended intrinsic maps. *ACM Trans. Graph.*, 30(4), July 2011. ISSN 0730-0301. doi: 10.1145/2010324.1964974.
- [35] Thomas N Kipf and Max Welling. Variational graph auto-encoders. *NIPS Workshop on Bayesian Deep Learning*, 2016.
- [36] Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations (ICLR)*, 2017.

- [37] Stefan Klus and Tuhin Sahai. A spectral assignment approach for the graph isomorphism problem. *ArXiv*, abs/1411.0969, 2014.
- [38] Oliver Koch, Nils M. Kriege, and Lina Humbeck. Chemical similarity and substructure searches. In *Encyclopedia of Bioinformatics and Computational Biology*, 2019.
- [39] Giorgos Kollias, Shahin Mohammadi, and Ananth Grama. Network similarity decomposition (nsd): A fast and scalable approach to network alignment. *IEEE Transactions on Knowledge and Data Engineering*, 24(12):2232–2243, 2012. doi: 10.1109/TKDE.2011.174.
- [40] Danai Koutra, Hanghang Tong, and David Lubensky. Big-align: Fast bipartite graph alignment. In *2013 IEEE 13th International Conference on Data Mining*, pages 389–398, 2013. doi: 10.1109/ICDM.2013.152.
- [41] Z. Löhner, E. Rodolà, M. M. Bronstein, D. Cremers, O. Burghard, L. Cosmo, A. Dieckmann, R. Klein, and Y. Sahillioğlu. Matching of deformable shapes with topological noise. In *Proceedings of the Eurographics 2016 Workshop on 3D Object Retrieval*, 3DOR '16, page 55–60, Goslar, DEU, 2016. Eurographics Association. ISBN 9783038680048.
- [42] Zorah Löhner, Matthias Vestner, Amit Boyarski, Or Litany, Ron Slossberg, Tal Remez, Emanuele Rodolà, Alexander M. Bronstein, Michael M. Bronstein, Ron Kimmel, and Daniel Cremers. Efficient deformable shape correspondence via kernel matching. *2017 International Conference on 3D Vision (3DV)*, pages 517–526, 2017.
- [43] M. Leordeanu and M. Hebert. A spectral technique for correspondence problems using pairwise constraints. In *Tenth IEEE International Conference on Computer Vision (ICCV'05) Volume 1*, volume 2, pages 1482–1489 Vol. 2, 2005. doi: 10.1109/ICCV.2005.20.
- [44] Marius Leordeanu. *Spectral Graph Matching, Learning, and Inference for Computer Vision*. PhD thesis, Carnegie Mellon University, Pittsburgh, PA, July 2009.
- [45] Qimai Li, Xiao-Ming Wu, Han Liu, Xiaotong Zhang, and Zhichao Guan. Label efficient semi-supervised learning via graph filtering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9582–9591, 2019.
- [46] Qimai Li, Xiaotong Zhang, Han Liu, Quanyu Dai, and Xiao-Ming Wu. Dimensionwise separable 2-d graph convolution for unsupervised and semi-supervised learning on graphs. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, KDD '21, page 953–963, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450383325. doi: 10.1145/3447548.3467413.
- [47] Yujia Li, Chenjie Gu, Thomas Dullien, Oriol Vinyals, and Pushmeet Kohli. Graph matching networks for learning the similarity of graph structured objects. In *ICML*, 2019.
- [48] Yaron Lipman and Thomas Funkhouser. Möbius voting for surface correspondence. *ACM Trans. Graph.*, 28(3), July 2009. ISSN 0730-0301. doi: 10.1145/1531326.1531378.
- [49] O. Litany, E. Rodolà, A. M. Bronstein, and M. M. Bronstein. Fully spectral partial shape matching. *Computer Graphics Forum*, 36(2):247–258, 2017. doi: <https://doi.org/10.1111/cgf.13123>.
- [50] Jingen Liu, Jiebo Luo, and Mubarak Shah. Recognizing realistic actions from videos “in the wild”. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1996–2003, 2009. doi: 10.1109/CVPR.2009.5206744.
- [51] Vince Lyzinski, Donniell E. Fishkind, and Carey E. Priebe. Seeded graph matching for correlated erdős-rényi graphs. *J. Mach. Learn. Res.*, 15(1):3513–3540, January 2014. ISSN 1532-4435.
- [52] Vince Lyzinski, Donniell E. Fishkind, Marcelo Fiori, Joshua T. Vogelstein, Carey E. Priebe, and Guillermo Sapiro. Graph matching: Relax at your own risk. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(1):60–73, 2016. doi: 10.1109/TPAMI.2015.2424894.
- [53] Riccardo Marin, Simone Melzi, Emanuele Rodolà, and Umberto Castellani. Farm: Functional automatic registration method for 3d human bodies. *Comput. Graph. Forum*, 39(1):160–173, 2020. URL <https://doi.org/10.1111/cgf.13751>.

- [54] Simone Melzi, Jing Ren, Emanuele Rodolà, Abhishek Sharma, Peter Wonka, and Maks Ovsjanikov. Zoomout: Spectral upsampling for efficient shape correspondence. *ACM Transactions on Graphics (TOG)*, 38(6):155, 2019.
- [55] Maks Ovsjanikov, Jian Sun, and Leo Guibas. Global intrinsic symmetries of shapes. *Comp. Graph. Forum*, 27(5):1341–1348, 2008.
- [56] Maks Ovsjanikov, Mirela Ben-Chen, Justin Solomon, Adrian Butscher, and Leonidas Guibas. Functional maps: A flexible representation of maps between shapes. *ACM Trans. Graph.*, 31(4), July 2012. ISSN 0730-0301. doi: 10.1145/2185520.2185526.
- [57] Gautam Pai, Jing Ren, Simone Melzi, Peter Wonka, and Maks Ovsjanikov. Fast sinkhorn filters: Using matrix scaling for non-rigid shape correspondence with functional maps. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 384–393, June 2021.
- [58] Jing Ren, Adrien Poulenard, Peter Wonka, and Maks Ovsjanikov. Continuous and orientation-preserving correspondences via functional maps. *ACM Trans. Graph.*, 37(6), dec 2018. ISSN 0730-0301. doi: 10.1145/3272127.3275040.
- [59] Jing Ren, Mikhail Panine, Peter Wonka, and Maks Ovsjanikov. Structured regularization of functional map computations. *Comput. Graph. Forum*, 38(5):39–53, 2019. doi: 10.1111/cgf.13788. URL <https://doi.org/10.1111/cgf.13788>.
- [60] E. Rodolà, L. Cosmo, M. M. Bronstein, A. Torsello, and D. Cremers. Partial functional correspondence. *Comput. Graph. Forum*, 36(1):222–236, January 2017. ISSN 0167-7055. doi: 10.1111/cgf.12797.
- [61] Emanuele Rodolà, Samuel Bulò, Thomas Windheuser, Matthias Vestner, and Daniel Cremers. Dense non-rigid shape correspondence using random forests. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pages 4177–4184, 2014. doi: 10.1109/CVPR.2014.532.
- [62] Yusuf Sahillioğlu and Yücel Yemez. Minimum-distortion isometric shape correspondence using em algorithm. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(11):2203–2215, 2012. doi: 10.1109/TPAMI.2012.26.
- [63] Aliaksei Sandryhaila and José M. F. Moura. Discrete signal processing on graphs. *IEEE Transactions on Signal Processing*, 61(7):1644–1656, 2013. doi: 10.1109/TSP.2013.2238935.
- [64] Gil Shamai and Ron Kimmel. Geodesic distance descriptors. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [65] Tianjia Shao, Wilmot Li, Kun Zhou, Weiwei Xu, Baining Guo, and Niloy J. Mitra. Interpreting concept sketches. *ACM Trans. Graph.*, 32(4), July 2013. ISSN 0730-0301. doi: 10.1145/2461912.2462003.
- [66] Roded Sharan and Trey Ideker. T. modeling cellular machinery through biological network comparison. *nat. biotechnol.* 24, 427-433. *Nature biotechnology*, 24:427–33, 05 2006. doi: 10.1038/nbt1196.
- [67] D. I. Shuman, S. K. Narang, P. Frossard, A. Ortega, and P. Vandergheynst. 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, 2013. doi: 10.1109/MSP.2012.2235192.
- [68] Rohit Singh, Jinbo Xu, and Bonnie Berger. Global alignment of multiple protein interaction networks with application to functional orthology detection. *Proceedings of the National Academy of Sciences of the United States of America*, 105:12763–8, 10 2008. doi: 10.1073/pnas.0806627105.
- [69] Justin Solomon, Andy Nguyen, Adrian Butscher, Mirela Ben-Chen, and Leonidas Guibas. Soft maps between surfaces. *Comput. Graph. Forum*, 31(5):1617–1626, August 2012. ISSN 0167-7055. doi: 10.1111/j.1467-8659.2012.03167.x.

- [70] Federico Tombari, Samuele Salti, and Luigi Di Stefano. Unique signatures of histograms for local surface description. In *European conference on computer vision*, pages 356–369. Springer, 2010.
- [71] Huynh Thanh Trung, Tong Van Vinh, Nguyen Thanh Tam, Hongzhi Yin, Matthias Weidlich, and Nguyen Quoc Viet Hung. Adaptive network alignment with unsupervised and multi-order convolutional networks. In *2020 IEEE 36th International Conference on Data Engineering (ICDE)*, pages 85–96, 2020. doi: 10.1109/ICDE48307.2020.00015.
- [72] S. Umeyama. An eigendecomposition approach to weighted graph matching problems. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 10(5):695–703, 1988. doi: 10.1109/34.6778.
- [73] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph Attention Networks. *International Conference on Learning Representations*, 2018. accepted as poster.
- [74] Mario Vento and Pasquale Foggia. Graph matching techniques for computer vision. *Graph-Based Methods in Computer Vision: Developments and Applications*, 1:1–41, 01 2012. doi: 10.4018/978-1-4666-1891-6.ch001.
- [75] Chun Wang, Shirui Pan, Ruiqi Hu, Guodong Long, Jing Jiang, and Chengqi Zhang. Attributed graph clustering: a deep attentional embedding approach. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence*, pages 3670–3676, 2019. doi: 10.24963/ijcai.2019/509.
- [76] Runzhong Wang, Junchi Yan, and Xiaokang Yang. Learning combinatorial embedding networks for deep graph matching. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, Oct 2019. doi: 10.1109/iccv.2019.00315.
- [77] Zhichun Wang, Qingsong Lv, Xiaohan Lan, and Yu Zhang. Cross-lingual knowledge graph alignment via graph convolutional networks. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 349–357. Brussels, Belgium, 2018. Association for Computational Linguistics. doi: 10.18653/v1/D18-1032.
- [78] Felix Wu, Amauri Souza, Tianyi Zhang, Christopher Fifty, Tao Yu, and Kilian Weinberger. Simplifying graph convolutional networks. In *Proceedings of the 36th International Conference on Machine Learning*, pages 6861–6871. PMLR, 2019.
- [79] Rui Xiang, Rongjie Lai, and Hongkai Zhao. Efficient and robust shape correspondence via sparsity-enforced quadratic assignment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9513–9522, 2020.
- [80] Rui Xiang, Rongjie Lai, and Hongkai Zhao. A dual iterative refinement method for non-rigid shape matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15930–15939, June 2021.
- [81] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? In *International Conference on Learning Representations*, 2019.
- [82] Lei Xu and Irwin King. A pca approach for fast retrieval of structural patterns in attributed graphs. *IEEE transactions on systems, man, and cybernetics. Part B, Cybernetics : a publication of the IEEE Systems, Man, and Cybernetics Society*, 31:812–7, 02 2001. doi: 10.1109/3477.956043.
- [83] Lyudmila Yartseva and Matthias Grossglauser. On the performance of percolation graph matching. In *COSN '13*, 2013.
- [84] LIREN YU, Jiaming Xu, and Xiaojun Lin. Graph matching with partially-correct seeds. *J. Mach. Learn. Res.*, 22:280:1–280:54, 2021.
- [85] Andrei Zanfir and Cristian Sminchisescu. Deep learning of graph matching. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2684–2693, 2018. doi: 10.1109/CVPR.2018.00284.

- [86] Jiawei Zhang and Philip S. Yu. Multiple anonymized social networks alignment. In *2015 IEEE International Conference on Data Mining*, pages 599–608, 2015. doi: 10.1109/ICDM.2015.114.
- [87] Xiaotong Zhang, Han Liu, Qimai Li, and Xiao-Ming Wu. Attributed graph clustering via adaptive graph convolution. In *IJCAI*, 2019.
- [88] Zhen Zhang and Wee Sun Lee. Deep graphical feature learning for the feature matching problem. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5086–5095, 2019. doi: 10.1109/ICCV.2019.00519.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? **[Yes]** see Sections 1, 3 4, 5 and Appendix D
 - (b) Did you describe the limitations of your work? **[Yes]** see Section 5.7
 - (c) Did you discuss any potential negative societal impacts of your work? **[Yes]** see Section 7
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? **[Yes]**
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? **[Yes]** see Sections 3, 2.2 and Appendix C
 - (b) Did you include complete proofs of all theoretical results? **[Yes]** see Sections listed in Question 2 (a) above.
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? **[Yes]** see code in supplementary material
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? **[Yes]** see Section 5.5 and Appendix D
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? **[N/A]**
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? **[Yes]** see Section 5.1
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? **[Yes]** see Section 5.5 and Appendix D.4
 - (b) Did you mention the license of the assets? **[N/A]**
 - (c) Did you include any new assets either in the supplemental material or as a URL? **[Yes]** see code in supplementary material
 - (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? **[Yes]** see Section 5.5 and Appendix D.4
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? **[N/A]**
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? **[N/A]**
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? **[N/A]**
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? **[N/A]**