
Solving Partial Assignment Problems using Random Clique Complexes

Charu Sharma¹ Deepak Nathani¹ Manohar Kaul¹

Abstract

We present an alternate formulation of the *partial assignment problem* as matching *random clique complexes*, that are higher-order analogues of random graphs, designed to provide a set of invariants that better detect higher-order structure. The proposed method creates random clique adjacency matrices for each k -skeleton of the random clique complexes and matches them, taking into account each point as the affine combination of its geometric neighborhood. We justify our solution theoretically, by analyzing the runtime and storage complexity of our algorithm along with the asymptotic behavior of the quadratic assignment problem (QAP) that is associated with the underlying *random clique adjacency matrices*. Experiments on both synthetic and real-world datasets, containing severe occlusions and distortions, provide insight into the accuracy, efficiency, and robustness of our approach. We outperform diverse matching algorithms by a significant margin.

1. Introduction

The *assignment problem* finds an assignment, or matching, between two finite sets U and V , each of cardinality n , such that the total cost of all matched pairs is minimized. The assignment problem can also be generalized to finding matchings between more than two sets. This is a fundamental problem in computer science and has been motivated by a wide gamut of research areas spanning diverse areas such as structural biology (Singer & Shkolnisky, 2011), protein structure comparisons in bioinformatics (Zaslavskiy et al., 2009), and computer vision (Conte et al., 2004). Computer vision especially boasts a broad range of applications that include object matching, image registration (Shen & Davatzikos, 2002), stereo matching (Goesele et al., 2007), shape matching (Pettersson et al., 2009; Berg et al., 2005),

¹Department of Computer Science & Engineering, Indian Institute of Technology Hyderabad, Hyderabad, India. Correspondence to: Charu Sharma <charusharma1991@gmail.com>, Deepak Nathani <deepakn1019@gmail.com>, Manohar Kaul <mkaul@iith.ac.in>.

Proceedings of the 35th International Conference on Machine Learning, Stockholm, Sweden, PMLR 80, 2018. Copyright 2018 by the author(s).




Figure 1. Matching cliques of two houses with 1-cliques (red), 2-cliques (green), 3-cliques (blue). Violet and yellow lines show the matchings of d -cliques where $d = 2, 3$ respectively.

structure from motion (SfM) (Szeliski, 2010), and object detection (Jiang et al., 2011), to name a few.

Various assignment approaches can broadly be classified as those that find a bijective assignment in the form of a permutation matrix by posing the problem as a *linear assignment problem (LP)* versus ones that solve a *quadratic assignment problem (QAP)* via *graph matching*, where each graph’s nodes represent the objects and the edges encode their corresponding distances; the goal of QAP then is to find node-wise correspondences between the graphs so that the overall discrepancy between their corresponding edge-wise counterparts is minimized and the overall *relational structure* is best preserved.

Partial assignment implies that only subsets of U and V can actually be assigned to each other successfully. This phenomenon is of particular interest to applications where either objects are absent due to incomplete observations, undergo deformations, and/or the objects in question cannot clearly be disambiguated because the objects in question along with their related objects are embedded in clutter. This variant of the assignment problem is widely accepted as a formidable challenge.

Although graph matching methods were found to be instrumental, they too perform poorly when faced with non-similar geometric transformations or transformations that produce degenerate triangulations. This is attributed to assigning weights to only node and edge assignments, while ignoring the interplay of higher-order connections/relations. For example, using triplet weights can alleviate the above

mentioned problem to a very large extent by defining a measure invariant to scale and other transformations (Chertok & Keller, 2010).

Motivated by the aforementioned observations and inspired by Kahle (Kahle, 2006)’s work on combinatorial topological models like the *random clique complex*, we focus our attention to matching higher-order components between two sets of points in the setting of some points *missing completely at random*. We pose our assignment problem as finding a matching between two sets of points, each represented as a random clique complex, which is a higher-order analogue of random graphs. Figure 1 illustrates such a matching of cliques of corresponding dimensionality, between two different scenes (taken from different camera angles) of the same house. Given an Erdős-Rényi (ER) graph, its *clique complex* is the simplicial complex with all complete subgraphs (i.e., cliques) as its *faces*. The Erdős-Rényi graph forms the 1-skeleton of the random clique complex, where the cliques have at most a dimension of 2, i.e., edges in the graph. The clique topology of a random adjacency matrix is analogous to its eigenvalue spectrum, as it provides a set of invariants that help detect structure (Giusti et al., 2015). This probabilistic and combinatorial framework of random clique complexes allows us to further study the assignment problem under various assumptions of the underlying distribution of the matrix entry distributions, its robustness to missing values, and its asymptotic behavior for large-scale cases.

Contributions: We present the following contributions.

1. To the best of our knowledge, our proposed approach is a first attempt to formulate *higher-order matching* between two sets of points, given partial or incomplete information, as a matching between two *random clique complexes*. We also propose an efficient matching algorithm and study both its time and storage complexity.
2. (i) We provide new bounds on the concentration inequality of eigenvalues of the QAP trace formulation for random symmetric matrices, (ii) we give tighter concentration inequality bounds on the largest eigenvalue for the Lawler QAP formulation on random matrices, in the context of *affinity matrices* that are used by some earlier works. Furthermore, we theoretically analyze and discuss the robustness of affinity-matrix based schemes to missing points, and (iii) we perform asymptotic analysis on the worst to best case ratio of a QAP solution for our higher-dimensional clique adjacency matrices in the *clique percolation regime* (Bollobás & Riordan, 2009), where the entries follow a *Poisson* distribution.
3. Finally, we present a comprehensive empirical study that compares our method’s matching accuracy to that

of a diverse set of matching approaches (Zhou & De la Torre, 2016; Zhou & De la Torre, 2013; Cho et al., 2010; Feizi et al., 2016; Leordeanu & Hebert, 2005; Cour et al., 2007; Pachauri et al., 2013; Gold & Rangarajan, 1996; Kuhn, 1955; Leordeanu et al., 2009; Zass & Shashua, 2008; Li et al., 2013; Duchenne et al., 2011). We conducted our experiments on both synthetic and well-known hard real-world datasets that span across affine/non-affine transformations, severe occlusions, and clutter. Our study reveals much better accuracy for the popular datasets against several of the state-of-the-art matching methods.

2. Matching Random Clique Complexes

We consider the problem of capturing higher-order feature groups among landmark points in an image by representing them as a *random clique complex* (RCC) and then using these RCCs to match two sets of groupings from two different images. We begin this section by describing the construction of a random clique complex, followed by our proposed method of matching two RCCs, and we finally analyze the runtime and storage complexity of our algorithm.

2.1. Structure of a Random Clique Complex

We begin with general definitions pertaining to the structure of simplicial complexes and then accordingly adapt these definitions to our domain of random graphs to build random clique complexes.

Let $G(n, p)$ be an Erdős-Rényi graph with a set of n vertices denoted by V , whose edges $\{v, v'\} \in \binom{V}{2}$, are i.i.d Bernoulli(p) distributed. Recall, that a k -clique in $G(n, p)$ is a *complete subgraph* that comprises of k vertices and $\binom{k}{2}$ edges. Here onwards, for ease of notation, we will denote $G(n, p)$ as G . Given any affinely independent set $V = \{v_i\}_{i=0}^k$ of $(k + 1)$ points in \mathbb{R}^n , the k -simplex $\sigma^{(k)}$ is the *convex hull* of V , i.e., it is the set of all points of the form $w_0v_0 + \dots + w_kv_k$, where $\sum_{i=0}^k w_i = 1$ and $w_i \geq 0$ for all i . If we imagine the vertices of G embedded generically in \mathbb{R}^n , then each $(k + 1)$ -clique consisting of $k + 1$ vertices is represented by a k -dimensional simplex $\sigma^{(k)}$ in our random clique complex. For example, a 2-clique (edge) and a 3-clique (triangle) in G is represented as $\sigma^{(1)}$ and $\sigma^{(2)}$, respectively.

Given $0 \leq i \leq k$, the i -th face f_i of $\sigma^{(k)}$ is the subspace of points that satisfy $w_i = 0$; it is the $(k - 1)$ -simplex $\sigma^{(k-1)}$ whose vertices are all those of $\sigma^{(k)}$, except the i -th vertex. In other words, when $\sigma^{(k)}$ is a clique of G , then all its subsets are also cliques and hence considered faces of $\sigma^{(k)}$. For example, a 3-clique (triangle) has three 2-cliques (edges) in it.

With the aforementioned definitions in mind, we define our

Algorithm 1 Matching Random Clique Complexes

Input: $\mathcal{X}(G) = \{G^{(k,l)}\}_{k=0}^h$ and $\mathcal{X}(G') = \{G'^{(k,l)}\}_{k=0}^h$
 1: **for** $k = h \dots 0$ **do**
 2: Let M, M' be the total number of $(k+1)$ -cliques in $G^{(k,l)}$ and $G'^{(k,l)}$, respectively
 3: $\mathcal{L} := \{c_i^{(k)}\}_{i=0}^{M-1}$ # list of barycenters
 4: **for** $i = 0 \dots M-1$ **do**
 5: $\mathcal{N}_i := \mathcal{N}_i \cup \{g_{(x,:)}^{(k,l)} \mid x = i, g_{(x,y)}^{(k,l)} \neq 0\}$
 6: $\mathcal{N} := \mathcal{N} \cup \{\mathcal{N}_i\}$ # clique neighborhoods
 7: **end for**
 8: **for** $i = 0 \dots M-1$ **do**
 9: $\alpha_i := [\alpha_1, \dots, \alpha_{M-1}]^T$
 10: $\alpha := \alpha \cup \{\alpha_i\}$ # affine weight vectors
 11: **end for**
 12: Repeat steps 3–11 on $G'^{(k,l)}$ for $\mathcal{L}', \mathcal{N}'$ and α' .
 13: Build cost matrix $C^{(k)}$ from weights vectors α, α'
 14: $X_k^* := \mathbf{Kuhn-Munkres}(G^{(k,l)}, G'^{(k,l)}, C^{(k)})$
 15: **end for**

Return: $\{X_0^*, \dots, X_h^*\}$ # set of permutation matrices

random clique complex $\mathcal{X}(G)$ as the set of all cliques in G such that $\mathcal{X}(G) = \{\sigma \in [n] \mid \sigma \text{ is a clique of } G\}$. We denote a set of $(k+1)$ -cliques as $\mathcal{X}_k(G)$. Additionally, $\mathcal{X}(G)$ also satisfies the following conditions of a simplicial complex: (i) Any face in $\mathcal{X}(G)$ is also a simplex in $\mathcal{X}(G)$ and (ii) the intersection of any two simplexes σ_i, σ_j is a face (lower dimensional clique) of both σ_i and σ_j .

The faces of $\sigma^{(k)}$ are *copies* of $\sigma^{(j)}$ for $j < k$, which are glued together inductively. The k -skeleton of $\mathcal{X}(G)$, for $k \in \mathbb{N}$, is defined as the following quotient space

$$\mathcal{X}^{(k)}(G) := \left(\mathcal{X}^{(k-1)}(G) \cup \coprod_{\sigma: \dim \sigma = k} \sigma^{(k)} \right) / \sim$$

where \sim is the equivalence relation that identifies faces of $\sigma^{(k)}$ to the corresponding faces of $\sigma \in \mathcal{X}^{(j)}(G)$ where $j < k$. Finally, $\mathcal{X}(G) = \cup_{k=0}^{\infty} \mathcal{X}^{(k)}(G)$.

k -skeleton as adjacency matrix: Given a random graph G and its k -skeleton $\mathcal{X}^{(k)}(G)$ that contains all its $(k+1)$ -cliques, we follow the idea from Bollobás et. al. (Bollobás & Riordan, 2009), to represent $\mathcal{X}^{(k)}(G)$ as an adjacency matrix $G^{(k,l)}$ whose vertex set is the set of all $(k+1)$ -cliques in G and in which two vertices (i.e., $(k+1)$ -cliques) are *adjacent* when they share a common face that has a minimum of l vertices, where $k \geq 1$ and $1 \leq l \leq k$. Such an adjacency matrix is built for each k -skeleton and therefore $\mathcal{X}(G)$ is expressed as a set of matrices $\{G^{(k,l)}\}_{k=0}^h$, where $(k+1)$ is the dimension of the cliques.

2.2. Problem Setup

The problem of matching random clique complexes each of dimension h , is the estimation of a set of optimal bijective maps of the form $\mathcal{M}_i : \mathcal{X}^{(i)}(G) \rightarrow \mathcal{X}^{(i)}(G')$, for all $i \leq h$, subject to assignment constraints. This can be formulated as a *constrained quadratic assignment problem*, which can later be relaxed to a linear programming optimization problem.

Given two h -dimensional random clique complexes $\mathcal{X}(G) = \{G^{(k,l)}\}_{k=0}^h$ and $\mathcal{X}(G') = \{G'^{(k,l)}\}_{k=0}^h$, let $X = \{X_0, \dots, X_h\} \in \Pi$ be a set of permutation matrices such that X_k encodes assignments/matchings from $G^{(k,l)}$ to $G'^{(k,l)}$. The combinatorial matching requires the optimal set of permutation matrices that best *align* $\mathcal{X}(G)$ and $\mathcal{X}(G')$. More formally, this can be expressed as the following constrained optimization problem

$$\begin{aligned} \operatorname{argmin}_{X_0, \dots, X_h} \quad & \sum_{k=0}^h \|G^{(k,l)} X_k - X_k G'^{(k,l)}\|_F^2 \\ \text{subject to} \quad & \forall k \leq h, \mathbb{1}^T X_k = \mathbb{1}, X_k^T \mathbb{1} = \mathbb{1} \end{aligned} \quad (1)$$

2.3. Our Algorithm

At a high level, our goal is to minimize $\|\mathcal{X}(G) - \mathcal{X}(G')\|_{\mathcal{C}}$, where \mathcal{C} is a combinatorial distance between two random clique complexes. Traditional metrics like Hausdorff distance are not suitable here because random clique complexes are combinatorial topological spaces. Recall that $\mathcal{X}(G)$ is comprised of a family of k -skeletons $\{\mathcal{X}^{(k)}(G)\}_{k=0}^h$, where each k -skeleton contains cliques whose dimension is at most $k+1$ and $\mathcal{X}^{(k)}(G)$ has a maximum dimension h . The solution of the optimization problem outlined in Equation (1) aims to find a set of permutation matrices $\{X_1, \dots, X_h\}$ that minimizes the overall number of *misalignments* between equi-dimensional faces of $\mathcal{X}(G)$ and $\mathcal{X}(G')$, i.e., cliques belonging to the corresponding k -skeletons, and thus producing the optimal least cost assignment between $\mathcal{X}(G)$ and $\mathcal{X}(G')$.

Algorithm 1 presents our method to solve the combinatorial optimization problem (Equation (1)). In decreasing order of clique dimensionality, for a fixed dimension k and given the adjacency matrices $G^{(k,l)}$ and $G'^{(k,l)}$ for k -skeletons $\mathcal{X}^{(k)}(G)$ and $\mathcal{X}^{(k)}(G')$, respectively. In every iteration, our objective is to solve $\operatorname{argmin}_{X_k} \|G^{(k,l)} X_k - X_k G'^{(k,l)}\|_F^2$ to find the optimal permutation X_k^* . We assume the barycenters of every clique is pre-computed (Step 3). Next, the *neighborhood* \mathcal{N}_i of the i -th clique is computed as the set of entries with 1s in the i -th row of $G^{(k,l)}$ (Step 5). We denote the collection of every clique's neighborhood as \mathcal{N} (Step 6). An important objective of our method is to capture the *geometric properties* of the neighborhood of every clique. We achieve this by characterizing the i -th clique's barycenter

$c_i^{(k)}$ as an *affine combination* of the barycenters (in all dimensions) associated with the cliques in its corresponding neighborhood \mathcal{N}_i . Given an arbitrary clique's barycenter $c_i^{(k)}$, let $\{x_1^{(k)}, \dots, x_n^{(k)}\}$ denote the barycenters of its n adjacent cliques. Then, $c_i^{(k)}$ expressed as $\sum_{i=1}^n \alpha_i x_i^{(k)}$ is an affine combination of the $x_i^{(k)}$ s, if $\sum_{i=1}^n \alpha_i = 1$, i.e., the weights α_i sum to 1. Among all possible affine representations of $c_i^{(k)}$ we chose to use *least squares* to guarantee minimal error under L2-norm, and furthermore it assigns non-zero weights to each of its adjacent clique barycenters, thereby capturing the local geometric properties in its neighborhood. The weight vector α_i is then calculated for each clique (Step 9) and α denotes a collection of such weight vectors (Step 10). Next, a cost matrix is built by computing the L2-norm distance between weight vectors α and α' (Step 13). Finally, the *Kuhn-Munkres* (Kuhn, 1955) algorithm is invoked with both the adjacency matrices and the cost matrix, which arrives at the optimal assignment (Step 14). At the end of all iterations, our method returns a set of optimal assignments for matches between each k -skeleton for every dimension below h and the algorithm terminates. We refer the reader to our supplementary section for a working example.

2.4. Complexity Analysis

To begin our analysis, we must first ascertain the dimensionality of $G^{(k,l)}$, which is governed by the total number of $(k+1)$ -cliques that exist in the underlying random graph G . It is important to note that there doesn't exist any closed form solution to counting the number of cliques of a given dimension in G .

We consider the distribution of a random variable $X_n(k)$ counting the number of $(k+1)$ -cliques in a realization of G . We show in Appendix A of our supplementary material that this count is upper bounded by $(en/k)^k$, where e is *Euler's number*. This can be expressed in asymptotic notation as $O(n^k)$. As dimensionality increases, there occurs an explosion in the number of cliques. Fortunately, $G^{(k,l)}$ is a sparse matrix and its *effective dimensionality* measured by the number of non-zero rows, i.e., the number of cliques with non-empty neighborhoods, is of order $O(\text{nnz}(G^{(k,l)}))$. Therefore, we set out to count the number of non-zero entries in $G^{(k,l)}$.

We use a seminal result by Bollobás (Bollobás & Riordan, 2009), where they identify a threshold probability for *percolation* of cliques in G for all fixed k and l , which is given by $p = \Theta\left(n^{\frac{-2}{k+l-1}}\right)$. Moreover, they proved that for p around this threshold, the number of cliques asymptotically converge to a *Poisson* distribution. Exceeding this threshold results in formation of *giant connected clique clusters*, which causes an explosion in the number of possible cliques.

Recall from our definition of $G^{(k,l)}$, that two cliques are adjacent if they share at least l vertices. In order to analyze this further, we imagine an entry in $G^{(k,l)}$ occurs when we can migrate a $(k+1)$ -clique from its original position to an adjacent clique by relocating exactly $(k+1-l)$ vertices and leaving the remaining l vertices intact. The expected number of such relocations is given by $\binom{k+1}{l} \binom{n}{k+1-l} p^{\binom{k+1}{2} - \binom{l}{2}}$, where the first term denotes the number of possible vertices in a $(k+1)$ -clique that can be chosen for relocation, the second term counts the number of new adjacent positions a clique can relocate to, and the final term decides the probability of relocations that are correct and acceptable. In our case, we define cliques to be adjacent to one another when they share at least $l = k$ vertices. This is done in order to keep the number of adjacent cliques to a manageable size during experiments. Setting $l = k$, gives $kn p^k$ expected relocations, which in turn estimates $\text{nnz}(G^{(k,l)})$.

Note that for every iteration in Algorithm 1, the dominating cost is that of running the Kuhn-Munkres matching algorithm in Step 14, which has a *cubic* cost in $\text{nnz}(G^{(k,l)})$. Let \mathcal{C}_{max} denote an upper bound on all the number of non-zero entries in $\{G^{(k,l)}\}_{i=0}^h$. Then, every iteration has a runtime $O(\mathcal{C}_{max}^3)$ and therefore after h iterations the final cost is $O(h\mathcal{C}_{max}^3)$. Observe that as the dimensionality of the cliques increases in every iteration, p^k decays very sharply and hence drastically reduces $\text{nnz}(G^{(k,l)})$, which in turn reduces the overall matching cost. Finally, the storage complexity can simply be given as $O(h\mathcal{C}_{max})$.

3. Theoretical Analysis of QAP

In this section, we present three related results in the context of matching random matrices, namely: (i) concentration inequality of eigenvalue bounds on the *QAP trace formulation* for random symmetric matrices, (ii) tighter concentration inequality of eigenvalue bounds on the *Lawler QAP formulation* on random symmetric matrices in the context of works that use *affinity matrices*, and (iii) provide an asymptotic analysis on the worst to best case ratio of a QAP for higher-dimensional clique adjacency matrices. For ease of notation, we will refer to the random clique adjacency matrices simply as A and B .

3.1. Eigenvalue Bounds of Trace QAP Formulation on Random Matrices

Let $A = (a_{vv'})$, $B = (b_{vv'}) \in \mathbb{R}^{n \times n}$ be random real-symmetric matrices. Let $X = (x_{ij}) \in \mathbb{R}^{n \times n}$ be a *permutation matrix*. Then, the trace formulation of a QAP is given by

$$\text{minimize } \text{Tr}(AXBX^T) \text{ s.t. } X \in \Pi_X$$

where Π_X is the set of permutation matrices.

Let $\lambda_1(A) \leq \lambda_2(A) \leq \dots \leq \lambda_n(A)$ and $\lambda_1(B) \geq \lambda_2(B) \geq \dots \geq \lambda_n(B)$ ¹ be the eigenvalues of A and B , respectively. Let the corresponding eigen-decompositions of matrices A and B , be given by $A = Q_A \Lambda_A Q_A^T$ and $B = Q_B \Lambda_B Q_B^T$, where $\Lambda_A = \text{diag}(\lambda_1(A), \dots, \lambda_n(A))$ and $\Lambda_B = \text{diag}(\lambda_1(B), \dots, \lambda_n(B))$ with their corresponding orthogonal eigenvector matrices Q_A and Q_B . Finke et. al. (Martello et al., 1987) gave the following eigenvalue bounds.

Theorem 1. *Let A and B be symmetric matrices. Then for all $X \in X_{\Pi}$,*

1. $\text{Tr}(AXBX^T) = \lambda(A)^T Q^{(X)} \lambda(B)$,
where $Q^{(X)} = \langle Q_A^{(i)}, X Q_B^{(j)} \rangle^2$ with vectors of eigenvalues given by $\lambda(A) = (\lambda_i(A))$ and $\lambda(B) = (\lambda_i(B))$. $Q_A^{(i)}$ and $Q_B^{(i)}$ denote the i -th eigenvectors of A and B , respectively;
2. $\mathcal{L} \leq \text{Tr}(AXBX^T) \leq \mathcal{U}$, where
 $\mathcal{L} = \sum_{i=1}^n \lambda_i(A) \lambda_i(B)$, and
 $\mathcal{U} = \sum_{i=1}^n \lambda_{n-i+1}(A) \lambda_i(B)$

It was also noticed by Finke (Martello et al., 1987) that these bounds can further be tightened by reducing the *spreads* of matrices A and B , where the spread of a matrix A , denoted by $\mathfrak{S}(A)$, is given by $\mathfrak{S}(A) = \max_i \lambda_i(A) - \min_i \lambda_i(A)$. There is no formula to compute the spread of a matrix directly, so Finke et. al. (Martello et al., 1987) suggested a *reduction* method to further sharpen the bound by replacing matrices A and B by *smaller spread* symmetric matrices \tilde{A} and \tilde{B} . The reductions are achieved as $\tilde{A} = A - M_A - M_A^T - \mathcal{D}_A$ and $\tilde{B} = B - M_B - M_B^T - \mathcal{D}_B$, where M_A, M_B are matrices with constant columns and $\mathcal{D}_A, \mathcal{D}_B$ are diagonal matrices, whose values are chosen appropriately in order to tighten the bounds on spreads $\mathfrak{S}(\tilde{A})$ and $\mathfrak{S}(\tilde{B})$.

Our bounds on Random Matrices: We propose new measure concentration inequalities on the spread of a random matrix, by redefining the spread in an alternate fashion that is more amenable to our analysis. Consider our reduced random symmetric matrix $\tilde{A} \in \mathbb{R}^{n \times n}$, with eigenvalues $\lambda_1(\tilde{A}) \leq \dots \leq \lambda_n(\tilde{A})$, we define the *gap* (spacing) between its consecutive eigenvalues as $\delta_i(\tilde{A}) := |\lambda_{i+1}(\tilde{A}) - \lambda_i(\tilde{A})|$ for $1 \leq i \leq n - 1$. Then, the spread $\mathfrak{S}(\tilde{A})$ for the reduced matrix \tilde{A} can be redefined as: $\mathfrak{S}(\tilde{A}) = \sum_{i=1}^n \delta_i(\tilde{A})$

We begin by upper bounding $\delta_i(\tilde{A})$ using the following lemma 1 (proof in supplementary notes).

Lemma 1. *Let $\|\cdot\|$ denote an algebraic matrix norm on a space of real $n \times n$ matrices \mathcal{M}_n , then for any $A \in \mathcal{M}_n$,*

$$\delta_i(A) \leq 2\|A\|$$

¹ The two sets of eigenvalues differ in ordering.

To the best of the author's knowledge there does not exist a known distribution of eigenvalue gaps for a symmetric random matrix. We now attempt to give concentration inequalities for the tail probabilities of the sum of eigenvalue gaps, i.e., the spread. For our i.i.d. random matrix $A \in \mathbb{R}^{n \times n}$, consider the sequence of independent eigenvalue gaps $\delta_1(A), \dots, \delta_n(A)$, where each $\delta_i(A)$ is upper bounded by $2\|A\|$, as shown in Lemma 1. Let us denote their sum as $\mathcal{S}_n(A) := \delta_1(A) + \dots + \delta_n(A)$. As $\delta_1(A), \dots, \delta_n(A)$ are independent scalar random variables with $\delta_i(A) \leq \|A\|$ a.s, with mean $\mu_i(A)$ and variance $\sigma_i^2(A)$. Then, using Chernoff bounds, for any $\epsilon > 0$, we have

$$\mathbb{P}(|\mathcal{S}_n(A) - \mu| \geq \epsilon\sigma) \leq K \max\left(e^{(-p\epsilon^2)}, e^{(-p\epsilon\sigma/2\|A\|)}\right)$$

for some absolute constants $K, p > 0$. The Chernoff inequality above, shows that $\mathcal{S}_n(A)$ is sharply concentrated in the range $n\mu + O(\sigma\sqrt{n})$, when ϵ is not too large.

3.2. Eigenvalue bounds on Lawler's QAP on Random Affinity Matrices

In literature, many graph matching algorithms use Lawler's QAP formulation. Recall, $A = (a_{ij}), B = (b_{uv}) \in \mathbb{R}^{n \times n}$. Let $\Omega(a_{i,j}, b_{u,v})$ denote the *pairwise affinity score* of assigning the (i, j) -th entry in A to the (u, v) -th entry in B , implying that node a_i is matched to node b_u and node a_j to node b_v , simultaneously. Then, the *affinity matrix* $\mathcal{A} \in \mathbb{R}^{n^2 \times n^2}$ is given by $\mathcal{A}[(i-1)n+u, (j-1)n+v] = \Omega(a_{i,j}, b_{u,v})$ and the optimal assignment to Lawler's QAP is the one that maximizes the sum total pairwise affinity scores. Leordeanu et. al. (Leordeanu & Hebert, 2005) show via a spectral relaxation that Lawler's WAP reduces to solving $w^* = \text{argmax}_w \frac{w^T \mathcal{A} w}{w^T w}$, $w \in \mathbb{R}^{n^2}$. This is solved by finding the leading eigenvalue $\lambda_1(\mathcal{A})$.

As illustrated in (Alon et al., 2002), we also use *Talagrand's concentration inequality* (Talagrand, 1995). We provide a tighter bound in the case of our affinity matrix using *Rayleigh's quotient*.

Theorem 2. *For a random affinity matrix $\mathcal{A} \in \mathbb{R}^{m \times m}$ and for a positive constant t , $\mathbb{P}[|\lambda_1(\mathcal{A}) - \mathcal{M}| \geq t] \leq 4e^{-t^2/8}$, where \mathcal{M} is the median of $\lambda_1(\mathcal{A})$. \square*

Discussion: We further investigate the robustness of affinity-matrix based graph matching solutions when dealing with *missing* or *incomplete* data. We show the sharpness of our result in Therem 2 on the affinity matrix, similar to (Alon et al., 2002), who analyze their results using *fat* matrices as an example. Consider an affinity matrix $\mathcal{A} = (a_{ij}) \in \mathbb{R}^{m \times m}$, whose entries are i.i.d. Bernoulli distributed. Simulating missing affinity scores due to missing edge assignments, we set $a_{ij} = 1$ with probability $1/4$ and $a_{ij} = 0$ with probability $3/4$. Notice that our random

affinity matrix represents the Erdős-Rényi graph $G(m, 1/4)$. As shown in (Alon et al., 2002), the median and expectation of $\lambda_1(A)$ differ by a constant factor. Let $G = (V, E)$ denote a general undirected graph, where the degree of each vertex $v \in V$ is given by $d_G : V \rightarrow \mathbb{Z}$. Then, the *average degree* of G is given by $\bar{d} = \sum_{v \in V} d_G(v)/|V|$ and its *maximum degree* is $\Delta = \max_{v \in V} d_G(v)$. It is then well known that $\bar{d} \leq \lambda_1(A) \leq \Delta$, i.e., the largest eigenvalue of a graph is squeezed between its average and maximum degree.

Let $|E|$ denote the total number of edges in $G(m, 1/4)$, then the *average degree* of $G(m, 1/4)$ is given by $2|E|/n$, where $|E| = \text{Bin}\left(\binom{m}{2}, 1/4\right)$ and the standard deviation of $|E|$ is $\sqrt{\binom{m}{2}(1/4)(3/4)} = \Theta(m)$. For large m , our binomial distribution converges to a normal distribution. Therefore, we calculate the probability for the total number of edges $|E|$ to deviate from its expectation by t standard deviations as $e^{-\Theta(t^2)}$. Furthermore, we know that if $|E|$ exceeds its expectation by $\Theta(tn)$, then the average degree \bar{d} must also correspondingly exceed its expectation by $\Theta(t)$. Therefore, the probability of the average degree \bar{d} exceeding its expectation by t standard deviations is *at least* $e^{-\Theta(t^2)}$. Given that $\bar{d} \leq \lambda_1(A)$, it follows that λ_1 exceeding its expectation is also lower bounded by the same $e^{-\Theta(t^2)}$. The bounds achieved are *tight up to a constant factor* in the exponent. Our experimental results on *Factorized Graph Matching (FGM)* by Zhou et. al. (Zhou & De la Torre, 2016) and *Re-weighted Random Walk Matching (RRWM)* (Cho et al., 2010) also support the finding that affinity matrix based matching solutions are more robust to missing edges due to occlusions in data.

3.3. Asymptotic Analysis of Higher-Order Clique Assignment

Following along the same lines as Finke et. al. (Martello et al., 1987), we study the asymptotic behavior of the worst to most optimal ratio and present it as the following theorem.

Theorem 3. *Given random clique adjacency matrices $A_n^{k,l}, B_n^{k,l} \sim \text{Pois}(\lambda)$ and their associated cost matrix $C_{vv'} \sim \text{Pois}(\lambda)$. We denote by $\lambda_e := \mathbb{E}(C_{vv'})$ and $\lambda_v := \mathbb{V}(C_{vv'})$ the expectation and variance of our Poisson distributed cost function. For $\epsilon > 0$ and $p = \Theta\left(n^{\frac{-2}{k+l-1}}\right)$, we have the following bound on the ratio of the worse to the best solution as*

$$\mathbb{P} \left\{ \frac{\max_{\pi \in \Pi} \sum_{vv'} C_{vv'}}{\min_{\pi \in \Pi} \sum_{vv'} C_{vv'}} \leq 1 + \epsilon \right\} \geq 1 - 2|\Pi| \exp \left(-2|S_\pi| \left(\frac{\epsilon' \sqrt{\lambda_v}}{\epsilon' + 2\lambda_v} \right)^2 \right) =: \psi(n, \epsilon)$$

where, $|\Pi| = n!$, $|S_\pi| = \binom{n+1}{2}$, $\lim_{n \rightarrow \infty} \psi(n, \epsilon) = 1$ \square

4. Experiments

Here, we study the robustness of various matching algorithms when affected by missing or incomplete information and transformations (both affine and non-affine) on synthetic and real-world datasets. For the sake of brevity, we report detailed dataset descriptions in our supplementary notes. The graph matching algorithms can broadly be classified based on their use of (i) *affinity-matrix*: FGM (Zhou & De la Torre, 2016; Zhou & De la Torre, 2013)², RRWM (Cho et al., 2010), (ii) *Eigenvalues*: EigenAlign (Feizi et al., 2016)³, SM (Leordeanu & Hebert, 2005), SMAC (Cour et al., 2007), PermSync (Pachauri et al., 2013)⁴, (iii) *LP relaxation*: GA (Gold & Rangarajan, 1996), Kuhn-Munkres (Kuhn, 1955), (iv) *Integer QAP*: IPFP (Leordeanu et al., 2009), (v) *Probabilistic matching*: PM (Zass & Shashua, 2008), (vi) *Higher-order matching given complete data*: Tensor (Duchenne et al., 2011)⁵, and (vii) *Geometric and Feature matching*: LAI-LP (Li et al., 2013)⁶.

4.1. Effect of Affine Transformations

Simulated Dataset: We perform affine transformations on *CMU House*, which is a sequence of N frames extracted from a video. More specifically, we uniformly sample frames (at 20% and 40%) and perform affine transformations on the selected frames to distort them. Figure 2 shows examples of affine transformations on house frame sequences. Table 1 shows the comparative error in matching for all the algorithms. We now describe each affine transformation as performance metrics in our experiments.

Rotation: Figure 2(b) shows a 180° rotated version of the original house frame (Figure 2(a)). Table 1 shows errors in matching when 20% of the frames are rotated by both 20° and 60°, respectively and when the same transformations are applied to 40% of the frames. As the percentage of transformed frames with greater degree increases, we note a substantial increase in error for other methods in comparison to our method's error increase.

Reflection: The reflected version of a house frame is shown in Figure 2(c). Table 1 shows that affinity-based approaches also performed equally well for reflection of house frame sequences.

Scaling: Resizing an image both horizontally and vertically scales the image as is shown in Figure 2(d). We fixed the scales to 0.5, 0.75, 1.25, and 1.5 randomly in both the directions in order to transform the images. Our method in Table 1 produces much better matchings than the other methods.

Shearing: We randomly apply shearing on house in one of

² FGM ³ EigenAlign ⁴ PermSync ⁵ Tensor ⁶ This algorithm serves as our *naive baseline* as it directly uses neighborhood properties of the underlying graph (LAI-LP).




Figure 2. (a) Original House Frame, (b)–(e) Four transformations on house frame: Rotation, Reflection, Scaling and Shear (green markers show true matching case with the original frame (a)).

Algorithms	20° Rotation		60° Rotation		Reflection		Scaling		Shear	
	20%	40%	20%	40%	20%	40%	20%	40%	20%	40%
OurMethod	0.01 ± 0.0	0.01 ± 0.0	0.05 ± 0.0	0.03 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.1 ± 0.0	0.2 ± 0.1	0.0 ± 0.0	0.0 ± 0.0
EigenAlign	63.97 ± 1.0	70.19 ± 1.0	65.39 ± 1.5	70.88 ± 1.9	62.5 ± 0.4	64.97 ± 0.2	62.37 ± 0.6	65.42 ± 1.1	61.04 ± 0.2	62.36 ± 0.7
FGM	3.6 ± 0.5	7.4 ± 0.5	18.0 ± 0.0	36.4 ± 0.5	0.0 ± 0.0	0.0 ± 0.0	2.20 ± 1.0	3.4 ± 0.5	0.0 ± 0.0	0.0 ± 0.0
LAI-LP	40.90 ± 0.9	43.07 ± 1.2	49.9 ± 1.2	61.16 ± 0.7	49.07 ± 0.9	59.41 ± 0.8	43.12 ± 5.5	45.91 ± 3.2	39.04 ± 0.7	38.27 ± 0.6
PermSync	13.36 ± 1.1	16.48 ± 0.4	25.48 ± 1.4	41.59 ± 0.7	26.24 ± 3.4	42.31 ± 5.1	12.15 ± 0.7	13.7 ± 0.6	10.19 ± 0.8	6.23 ± 3.2
RRWM	2.0 ± 0.0	4.0 ± 0.0	14.0 ± 0.0	27.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	5.0 ± 2.0	10.0 ± 1.5	0.0 ± 0.0	0.0 ± 0.0
Tensor	6.88 ± 0.0	14.39 ± 0.8	19.24 ± 0.6	37.95 ± 0.8	17.1 ± 0.5	34.69 ± 0.4	2.91 ± 0.6	6.19 ± 0.7	0.0 ± 0.0	0.0 ± 0.0
IPFP	6.8 ± 1.5	11.6 ± 0.5	18.2 ± 0.5	35.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	7.4 ± 2.0	15.6 ± 1.5	0.8 ± 0.5	0.6 ± 0.5
PM	43.2 ± 1.0	49.4 ± 1.5	46.4 ± 1.0	56.8 ± 0.5	37.0 ± 0.0	37.0 ± 0.0	45.6 ± 0.5	53.4 ± 0.5	37.8 ± 0.5	38.4 ± 0.5
SMAC	12.6 ± 0.5	20.0 ± 0.0	19.0 ± 0.0	33.0 ± 0.0	4.0 ± 0.0	4.0 ± 0.0	13.8 ± 2.5	23.4 ± 0.5	4.0 ± 0.0	4.2 ± 0.5
SM	32.6 ± 0.5	39.4 ± 1.5	37.8 ± 1.0	49.6 ± 1.0	25.0 ± 0.0	25.0 ± 0.0	33.0 ± 1.5	40.4 ± 0.5	25.2 ± 0.5	24.8 ± 0.5
GA	34.0 ± 1.5	37.0 ± 1.0	38.4 ± 1.0	47.0 ± 0.0	30.0 ± 0.0	30.0 ± 0.0	37.6 ± 1.0	45.4 ± 0.5	29.0 ± 0.0	28.4 ± 0.5
Munkres	35.25 ± 1.5	36.47 ± 0.9	44.75 ± 0.8	55.13 ± 1.3	46.05 ± 1.2	57.25 ± 2.2	37.32 ± 1.5	37.81 ± 1.8	32.84 ± 0.3	31.27 ± 0.8

Table 1. Error (%) of transformation on CMU House: inserted 20% and 40% impurity in CMU House frame sequence randomly for rotation (20°, 60°), reflection, scaling and shear. Minimum error (%) is shown in bold. Matching is computed for 111 frames from the 1st frame to the other 110 frames. Our method shows best performance among all the methods.

the directions with shear factor 0.5 (shown in Figure 2(e)) and measured the performance shown in Table 1. In addition to our method, we find that affinity-based algorithms also produce robust matchings.

4.2. Effect of Incomplete and Occluded Landmarks

To understand the effect of occlusions, we took two real-world datasets, i.e., *Books* and *Building* (Pachauri et al., 2013) with severe occlusions which are scenes of the same 3D object taken from arbitrary camera angles. These datasets have widely been used in *Structure from Motion (SfM)* problems and are known to be difficult for matching. Focusing our attention to the last two columns of Table 2, it is evident that our method gives the best results. Figure 3 shows the Books dataset where books are placed on a table in various orientations with varying levels of occlusion, along with two sample matchings between different pairs of images. Note that in Figure 3, when a corresponding matching clique is not found in the other image, a match isn’t forced but rather there is no match reported, which doesn’t degrade the matching accuracy. Matching as many random cliques, in order of decreasing dimensionality, as possible, manifests itself as an advantage over existing methods, especially when dealing with clutter and/or occlusions. **Simulating missing points:** In order to gain a deeper insight into the behavior of all the matching algorithms, we omit 2, 4, 6, 8, and 10 (6.66%, 13.33%, 20%, 26.66%, and



Figure 3. Two instances of matchings in Books dataset which is severely occluded. Yellow/green lines show correct/incorrect matches and isolated points show no matches.

33.33%) points out of total House landmark points (i.e., 30 points) from 40% (Figure 4) of frame sequences randomly. In general, all algorithms show an increase in error as more points are removed, but our method has a less gradual increase, while eigenvalue related methods show a rather steep increase in error. Our method is comparable to FGM and RRWM, but the gap in error increases with more missing points. We also observed that FGM incurs the longest runtime for matching in this scenario.

Methods	Car	Bike	Butterfly	Magazine	Building	Book
OurMethod	4.14 ± 2.45/ 7.13	3.15 ± 0.32/ 6.96	3.89 ± 0.23/ 14.76	0.48 ± 0.02/ 43.99	4.17 ± 0.32/ 12.65	22.20 ± 1.16/ 14.86
EigenAlign	60.68 ± 0.29/ 19.37	57.44 ± 0.37/ 19.32	66.57 ± 0.0/ 26.13	43.23 ± 0.0/ 93.60	90.51 ± 0.0/ 2.64	98.41 ± 0.0/ 8.29
FGM	55.51 ± 0.0/ 1793.9	48.17 ± 0.0/ 2013.7	16.12 ± 0.0/ 674.94	0.0 ± 0.0/ 777.55	74.87 ± 0.05/ 2530.5	97.54 ± 0.01/ 4293.9
LAI-LP	73.06 ± 0.23/ 152.47	42.00 ± 0.24/ 154.15	49.54 ± 0.0/ 161.06	88.73 ± 0.1/ 184.153	87.98 ± 0.0/ 33.06	96.38 ± 0.0/ 14.71
PermSync	10.63 ± 0.0/ 0.45	8.90 ± 0.0/ 0.46	46.93 ± 0.0/ 0.43	79.88 ± 0.0/ 1.08	64.00 ± 0.0/ 0.22	70.00 ± 0.0/ 0.48
RRWM	60.91 ± 0.0/ 4.96	54.53 ± 0.0/ 4.83	30.99 ± 0.0/ 8.53	1.98 ± 0.0/ 18.09	72.87 ± 0.01/ 7.98	87.04 ± 0.0/ 21.84
Tensor	24.37 ± 0.9/ 93.36	15.07 ± 1.0/ 93.97	1.07 ± 0.17/ 107.93	0.0 ± 0.0/ 182.07	43.24 ± 2.98/ 40.21	32.35 ± 0.15/ 40.41
IPFP	65.13 ± 0.0/ 6.35	60.81 ± 0.0/ 6.28	40.90 ± 0.0/ 8.43	3.94 ± 0.0/ 12.31	76.19 ± 0.0/ 4.65	87.74 ± 0.0/ 8.90
PM	74.63 ± 0.0/ 7.07	71.93 ± 0.0/ 4.90	70.27 ± 0.0/ 0.94	48.82 ± 0.0/ 1.69	83.79 ± 0.02/ 2.98	91.43 ± 0.24/ 0.44
SMAC	70.00 ± 0.0/ 5.75	67.36 ± 0.0/ 5.52	50.53 ± 0.0/ 4.15	5.52 ± 0.0/ 6.90	78.56 ± 0.22/ 1.94	87.88 ± 0.11/ 3.42
SM	68.54 ± 0.0/ 3.34	67.18 ± 0.0/ 3.47	65.96 ± 0.0/ 3.32	34.16 ± 0.0/ 4.93	80.27 ± 0.07/ 1.88	88.66 ± 0.09/ 2.10
GA	65.06 ± 0.0/ 4.53	64.60 ± 0.0/ 4.61	61.58 ± 0.0/ 4.08	31.62 ± 0.0/ 5.83	77.20 ± 0.27/ 3.51	87.02 ± 0.16/ 32.28
Munkres	33.71 ± 0.0/ 1.52	29.99 ± 0.0/ 1.49	51.87 ± 0.0/ 1.39	79.69 ± 0.0/ 2.45	74.00 ± 0.0/ 0.73	92.00 ± 0.0/ 1.16

Table 2. Car, Motorbike (Cho et al., 2013), Butterfly, Magazine (Jiang et al., 2011), Building and Books error (%) for pairwise matchings. Computation time (in seconds) is mentioned after the "/" in the above Table.

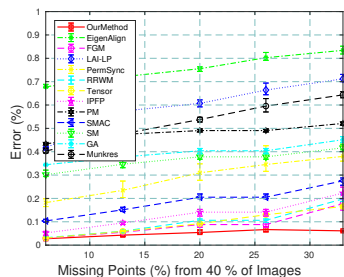


Figure 4. Error (%) in matching when varying the number of missing landmarks in 40% of the images in the frame sequence.




Figure 5. Error (%) in matching by various methods with different frame separation level for *CMU Hotel* (on left) and *Horse Shear* (on right).

4.3. Effect of Frame Separation

Here, we pick two frames from a video for matching and vary the separation in their frame sequence number. The farther apart two frames are the more *pronounced* is the effect we seek between the frame images. For example, as the frame separation increases, *CMU Hotel* undergoes a more severe 3D rotation, while *Horse-Shear* (Caetano et al., 2009) undergoes a larger degree of shear. We set $p = 0.7$ and $k = 7$ as nearest neighbors to get the correct matchings. In Figure 5, in both the left and right plots we notice that most methods show a very sharp rise in error, while our method is quite stable and reports a 0% error. We observe that the naive baseline, LAI-LP also does well and doesn't exhibit steep changes in error with larger frame separation.

Experimental Summary: In general, we find that the *affinity matrix* based methods like FGM and RRWM are more

robust to affine transformations than other competing algorithms. Our method performs the best as the weight vectors in our algorithm effectively capture even the *higher-order geometric properties* of the neighborhood and nearly preserves them under affine transformations. The naive baseline, i.e., LAI-LP, does not perform as well because it also has a feature-based component like SIFT which is known to fail on some affine transformations.

5. Conclusion

To the best of our knowledge, we have presented the first approach towards partial higher-order matching by initially capturing higher-order structure as *random clique complexes* and then proposing a corresponding matching algorithm. From a theoretical point of view, we studied matching as a QAP on random clique adjacency matrices that represented the k -skeleton of our random clique complexes and gave bounds on the concentration inequality of the spread of its eigenvalues. We also improved bounds on the largest eigenvalue of the Lawler QAP formulation, used by affinity-matrix based approaches. We discussed the robustness of such approaches to missing points and also showed the sharpness of our result. Furthermore, inspired by Finke et al. (Martello et al., 1987) we studied the asymptotic behavior of our higher dimensional clique adjacency matrices. A more detailed investigation of the distribution of eigenvalue gaps for such random matrices with Poisson distributed entries is left for future work.

From an empirical perspective, we compared the matching accuracies of diverse algorithms on both synthetic and real-world datasets that were known to have severe occlusions and distortions, thus posing a daunting challenge to matching algorithms. We argue that our experiments show strong evidence that our approach outperforms all the state-of-the-art matching methods on a diverse range of datasets.

Acknowledgements

We thank our colleagues from the Mathematics Dept. at IIT-H (Sukumar Daniel, Narasimha Kumar, and Bhakti B. Manna) for their insight and expertise. We would also like to thank all the reviewers for their feedback and suggestions. We are grateful to the authors of (Zhou & De la Torre, 2016; Zhou & De la Torre, 2013; Cho et al., 2010; Feizi et al., 2016; Pachauri et al., 2013; Li et al., 2013; Duchenne et al., 2011) for providing their source codes and datasets.

References

- Alon, N., Krivelevich, M., and Vu, V. H. On the concentration of eigenvalues of random symmetric matrices. *Israel Journal of Mathematics*, 131(1):259–267, Dec 2002.
- Berg, A. C., Berg, T. L., and Malik, J. Shape matching and object recognition using low distortion correspondences. In *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*, volume 1, pp. 26–33. IEEE, 2005.
- Bollobás, B. and Riordan, O. Clique percolation. *Random Struct. Algorithms*, 35(3):294–322, 2009.
- Caetano, T. S., McAuley, J. J., Cheng, L., Le, Q. V., and Smola, A. J. Learning graph matching. *IEEE transactions on pattern analysis and machine intelligence*, 31(6):1048–1058, 2009.
- Chertok, M. and Keller, Y. Efficient high order matching. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(12):2205–2215, Dec 2010.
- Cho, M., Lee, J., and Lee, K. M. Reweighted random walks for graph matching. In *European conference on Computer vision*, pp. 492–505. Springer, 2010.
- Cho, M., Alahari, K., and Ponce, J. Learning graphs to match. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 25–32, 2013.
- Conte, D., Foggia, P., Sansone, C., and Vento, M. Thirty years of graph matching in pattern recognition. *International journal of pattern recognition and artificial intelligence*, 18(03):265–298, 2004.
- Cour, T., Srinivasan, P., and Shi, J. Balanced graph matching. In *Advances in Neural Information Processing Systems*, pp. 313–320, 2007.
- Duchenne, O., Bach, F., Kweon, I.-S., and Ponce, J. A tensor-based algorithm for high-order graph matching. *IEEE transactions on pattern analysis and machine intelligence*, 33(12):2383–2395, 2011.
- Feizi, S., Quon, G., Recamonde-Mendoza, M., Médard, M., Kellis, M., and Jadbabaie, A. Spectral alignment of networks. *arXiv preprint arXiv:1602.04181*, 2016.
- Giusti, C., Pastalkova, E., Curto, C., and Itskov, V. Clique topology reveals intrinsic geometric structure in neural correlations. *Proc. Nat. Acad. Sci. USA*, 112(44):13455–13460, 2015.
- Goesele, M., Snaveley, N., Curless, B., Hoppe, H., and Seitz, S. M. Multi-view stereo for community photo collections. In *Computer Vision, 2007. ICCV 2007. IEEE 11th International Conference on*, pp. 1–8. IEEE, 2007.
- Gold, S. and Rangarajan, A. A graduated assignment algorithm for graph matching. *IEEE Transactions on pattern analysis and machine intelligence*, 18(4):377–388, 1996.
- Jiang, H., Stella, X. Y., and Martin, D. R. Linear scale and rotation invariant matching. *IEEE transactions on pattern analysis and machine intelligence*, 33(7):1339–1355, 2011.
- Kahle, M. Topology of random clique complexes, 2006. URL <http://arxiv.org/abs/math.CO/0605536>.
- Kuhn, H. W. The hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2:83–97, 1955.
- Leordeanu, M. and Hebert, M. A spectral technique for correspondence problems using pairwise constraints. In *Proceedings of the Tenth IEEE International Conference on Computer Vision - Volume 2, ICCV '05*, pp. 1482–1489, 2005.
- Leordeanu, M., Hebert, M., and Sukthankar, R. An integer projected fixed point method for graph matching and map inference. In *Advances in neural information processing systems*, pp. 1114–1122, 2009.
- Li, H., Huang, X., and He, L. Object matching using a locally affine invariant and linear programming techniques. *IEEE transactions on pattern analysis and machine intelligence*, 35(2):411–424, 2013.
- Martello, S., Minoux, M., Ribeiro, C., and Laporte, G. *Surveys in Combinatorial Optimization, Volume 31*. jan 1987.
- Pachauri, D., Kondor, R., and Singh, V. Solving the multi-way matching problem by permutation synchronization. In *Advances in neural information processing systems*, pp. 1860–1868, 2013.
- Petterson, J., Yu, J., McAuley, J. J., and Caetano, T. S. Exponential family graph matching and ranking. In *Advances in Neural Information Processing Systems*, pp. 1455–1463, 2009.

- Shen, D. and Davatzikos, C. Hammer: hierarchical attribute matching mechanism for elastic registration. *IEEE transactions on medical imaging*, 21(11):1421–1439, 2002.
- Singer, A. and Shkolnisky, Y. Three-dimensional structure determination from common lines in cryo-em by eigenvectors and semidefinite programming. *SIAM journal on imaging sciences*, 4(2):543–572, 2011.
- Szeliski, R. *Computer vision: algorithms and applications*. Springer Science & Business Media, 2010.
- Talagrand, M. Concentration of measure and isoperimetric inequalities in product spaces. *Publications Mathématiques de l’Institut des Hautes Études Scientifiques*, 81(1): 73–205, 1995.
- Zaslavskiy, M., Bach, F., and Vert, J.-P. Global alignment of protein–protein interaction networks by graph matching methods. *Bioinformatics*, 25(12):i259–1267, 2009.
- Zass, R. and Shashua, A. Probabilistic graph and hypergraph matching. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, pp. 1–8. IEEE, 2008.
- Zhou, F. and De la Torre, F. Deformable graph matching. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2922–2929, 2013.
- Zhou, F. and De la Torre, F. Factorized graph matching. *Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 38(9):1774–1789, 2016.