

Tight Relaxation of Quadratic Matching

Itay Kezurer[†] Shahar Z. Kovalsky[†] Ronen Basri Yaron Lipman

Weizmann Institute of Science




Figure 1: Consistent Collection Matching. Results of the proposed one-stage procedure for finding consistent correspondences between shapes in a collection showing strong variability and non-rigid deformations.

Abstract

Establishing point correspondences between shapes is extremely challenging as it involves both finding sets of semantically persistent feature points, as well as their combinatorial matching. We focus on the latter and consider the Quadratic Assignment Matching (QAM) model. We suggest a novel convex relaxation for this NP-hard problem that builds upon a rank-one reformulation of the problem in a higher dimension, followed by relaxation into a semidefinite program (SDP). Our method is shown to be a certain hybrid of the popular spectral and doubly-stochastic relaxations of QAM and in particular we prove that it is tighter than both.

Experimental evaluation shows that the proposed relaxation is extremely tight: in the majority of our experiments it achieved the certified global optimum solution for the problem, while other relaxations tend to produce sub-optimal solutions. This, however, comes at the price of solving an SDP in a higher dimension.

Our approach is further generalized to the problem of Consistent Collection Matching (CCM), where we solve the QAM on a collection of shapes while simultaneously incorporating a global consistency constraint. Lastly, we demonstrate an application to metric learning of collections of shapes.

1. Introduction

Establishing correspondences between shapes is a fundamental problem in computer graphics and related fields. It is a key ingredient in algorithms aiming at shape analysis and understanding.

Many different approaches have been suggested for both the formulation and computation of the extremely challenging shape correspondence problem. One common approach aims at finding a set of meaningful pairs of corresponding

points between two shapes, which in turn poses two main difficulties: (1) finding good candidate points on general shapes; and (2) the combinatorial problem of matching these candidates. We focus on the latter.

Given a set of feature points, a popular model for formulating the point correspondence problem is the Quadratic Assignment Matching (QAM) [BBM05, BBK06, FS06],

$$\max_{X \in \Pi_n} [X]^T W [X]$$

where $[X] \in \mathbb{R}^{n^2 \times 1}$ denotes the column-stack vector of X , and Π_n denotes the set of $n \times n$ permutation matrices.

[†] equal contributors

Notably, this problem is NP-hard and is considered to be large scale for $n \geq 16$ and generally intractable for $n \geq 30$ [ZKRW98, LdABN*07]. As such, different approaches have been suggested in computer graphics and computer vision to produce tractable algorithms for approximating its solution. Among these, the *spectral* and *doubly-stochastic* relaxations have stood out and became standard [LH05, Mem07, SLZ*13, FHJB13, ABK14].

In this paper we suggest a novel convex relaxation to the QAM problem that is *provably better* than both spectral and doubly-stochastic relaxations. It is achieved by a rank-one reformulation of the problem in a higher dimension, followed by relaxation into a semidefinite program (SDP).

Experiments show that the proposed convex relaxation usually achieves an integral solution and therefore globally optimal, while other relaxations often tend to produce sub-optimal solutions. On the other hand, since our approach requires solving an SDP in a rather high dimension, computation times are significantly higher than for simpler relaxation alternatives. Nevertheless, the method is suitable for computing coarse sets of point correspondences in reasonable times, often with a certificate of global optimality.

Going beyond pairwise matching, we generalize our framework to the problem of Consistent Collection Matching (CCM). Thereby, computing a point correspondences across an entire collection of shapes in a *single-step* algorithm, which simultaneously enforces a global consistency criterion. This is in contrast to previous techniques that tackle the CCM problem in two steps: first, independently matching many or all pairs of shapes in the collection, and second, combining all pairwise information in a globally consistent manner [NBCW*11, SW12, KLM*12, HG13, HWG14]. The power of the proposed CCM framework is demonstrated in Figure 1, where coarse correspondences are established over a collection of octopus shapes taken from the SHREC dataset [GBP07]; the correspondences, illustrated by colored spheres, further induce a continuous map between shapes, visualized by coloring their surfaces; see Section 5 for additional details. Lastly, we demonstrate an application to metric learning, whereby the CCM framework naturally induces a metric on the shape space.

2. Previous work

Shape correspondences and Quadratic Assignment Matching. The problem of shape correspondence in the context of computer graphics is surveyed in [vKZHC01]. Several studies have formulated the problem of establishing correspondences as the optimization of a well-defined quadratic objective. [BBK06, BBK*10] use a Gromov-Hausdorff framework to define the notion of partial similarity between two objects. This framework is also used by [Mem07] for shape comparison; he further describes its relation with the quadratic assignment problem and proposes a

sequence of constrained convex problems for its approximation. [BBM05] approximate the solution of QAM by solving a linear assignment problem, followed by local refinement. [LH05] rely on the Perron-Frobenius theorem to propose a spectral relaxation, followed by a greedy approach for enforcing integrality and mapping constraints. This spectral relaxation is used in [FHJB13] for matching graphs representing surfaces, derived from the Morse-Smale complex. [DBKP11] generalize the spectral approach to higher-order relations between tuples of points, represented by tensors. The closely related problem of graph matching was formulated as a QAM by [ABK14]; they characterize instances for which its doubly-stochastic relaxation is exact. The problem of establishing partial correspondences has been discussed in [FS06]; a priority-driven search approach has been proposed for large scale shape retrieval.

Quadratic Assignment. The QAM is an instance of the quadratic assignment problem (QAP) which is covered in a large body of literature, see [LdABN*07] for a survey. Several semidefinite programming relaxations have been suggested in this context. Most related is the work of [ZKRW98], which also present a certain semidefinite programming relaxation to a lifted reformulation of the problem; their formulation, however, is different: it does not include the properties we use for proving the relation to standard relaxations, nor does it address the case of sub-permutations.

3. Problem statement

The goal of this paper is to introduce an approach for solving the *Quadratic Assignment Matching* (QAM). The input of the problem is a pair of shapes \mathcal{M}^1 and \mathcal{M}^2 . Each shape $\mathcal{M}^i = \{\mathcal{P}^i, d^i\}$ consists of a *given* set of n points $\mathcal{P}^i = \{\mathbf{p}_1^i, \dots, \mathbf{p}_n^i\} \subset \mathbb{R}^d$, and a distance function $d(\mathbf{p}_q^i, \mathbf{p}_s^i)$ defined between pairs of these points.

Our goal is to establish good partial matching of order k between the shapes \mathcal{M}^1 and \mathcal{M}^2 . We model partial matchings, or injective mappings of order k ($0 \leq k \leq n$), via the set Π_n^k of $n \times n$ sub-permutations of rank k . That is, binary matrices with k ones, and at most a single one in each row and column. The inset shows an example of a matrix in Π_6^4 . The matching between \mathcal{M}^1 and \mathcal{M}^2 is represented by a sub-permutation matrix $X \in \Pi_n^k$, wherein \mathbf{p}_q^1 is matched to \mathbf{p}_r^2 if and only if $X_{qr} = 1$.

For the case of partial matching, the QAM problem takes the following matrix form

$$\max_X [X]^T W [X] \quad (1a)$$

$$\text{s.t. } X \in \Pi_n^k \quad (1b)$$

where $W \in \mathbb{R}^{n^2 \times n^2}$ defines a quadratic energy measuring the



gain in matching \mathcal{M}_1 and \mathcal{M}_2 by a sub-permutation X . As before, $[X] \in \mathbb{R}^{n^2 \times 1}$ denotes the column-stack vector of X .

A canonical instantiation of the QAM framework can be obtained by choosing the *isometric matching functional*, which intuitively aims to match pairs of points that preserve distances [LH05, BBM05, BBK06, FHJB13]. Formally, we set d^i to be the Euclidean distance (or geodesic in case of non-flat geometries) and the cost matrices W to decay proportionally to their differences, *i.e.*,

$$d(\mathbf{p}_q^i, \mathbf{p}_s^i) = \|\mathbf{p}_q^i - \mathbf{p}_s^i\|, \quad (2a)$$

$$W_{qrst} = e^{-|d(\mathbf{p}_q^1, \mathbf{p}_s^1) - d(\mathbf{p}_r^2, \mathbf{p}_t^2)|^2 / \sigma^2}. \quad (2b)$$

The QAM problem as formulated in (1) is generally non-convex, both in its objective and constraint. In fact, it is as hard as the quadratic assignment problem (QAP) and therefore NP-hard. Nonetheless, we propose an effective approach for its approximation.

4. Quadratic assignment matching relaxation

In this section we propose a novel SDP relaxation for the QAM problem (1). It is shown to outperform the popular spectral and doubly-stochastic relaxations for the QAM problem (*e.g.*, [LH05, ABK14]). In fact, the proposed formulation can be viewed as a hybrid of these two standard relaxations, which further provides a unifying convex interpretation of these techniques. See Figure 2 for an illustration.

We develop our convex relaxation to (1) in a few stages: first, we reformulate the problem as an equivalent linear problem subject to rank-one constraints, which is then relaxed into a semidefinite program. Then, we discuss its relation to standard QAM relaxations and show it can be further tightened.

4.1. Lifting and rank-one reformulation

Problem (1) is the maximization of a quadratic objective, *not necessarily concave*, over the set of $n \times n$ sub-permutation of order k . In this section, we show it can be *lifted* to the space of $n^2 \times n^2$ matrices. It can then be equivalently reformulated as the problem of maximizing a linear objective over the set of rank-one matrices. In subsequent section we further exploit its structure and propose an effective algorithm for its approximation.

The lifting is done by observing that (1) can equivalently be rewritten by introducing an auxiliary $n^2 \times n^2$ matrix variable Y that obeys $Y = [X][X]^T$. Problem (1) then becomes

$$\max_{X, Y} \quad \text{tr}(WY) \quad (3a)$$

$$\text{s.t.} \quad Y = [X][X]^T \quad (3b)$$

$$X \in \Pi_n^k \quad (3c)$$




Figure 2: Illustration of QAM relaxation hierarchy. The non-convex QAM problem (1) is illustrated as a star at the center. The intersection of the standard convex doubly-stochastic and non-convex spectral relaxations can be characterized as a convex SDP (6). Additional constraints yield the tight QAM-SDP relaxation (7).

With the lifting Y , the non-convex permutation constraint (3c) can be replaced by a convex constraint, as the following proposition asserts,

Proposition 1 Suppose $Y = [X][X]^T$. Then $X \in \Pi_n^k$ if and only if the following convex conditions hold,

$$X \in \text{conv} \Pi_n^k \quad (4a)$$

$$\text{tr} Y = k \quad (4b)$$

Here $\text{conv} \Pi_n^k$ is the convex-hull of the sub-permutations Π_n^k , characterized by the following linear inequalities and equality [DM58],

$$X \geq 0 \quad (5a)$$

$$X \mathbf{1} \leq \mathbf{1} \quad (5b)$$

$$X^T \mathbf{1} \leq \mathbf{1} \quad (5c)$$

$$\mathbf{1}^T X \mathbf{1} = k, \quad (5d)$$

where $\mathbf{1} \in \mathbb{R}^{n^2 \times 1}$ denotes the vector of all ones.

Proof If $X \in \Pi_n^k$ one can verify that (4a) and (4b) hold by substituting $Y = [X][X]^T$. In the other direction, assume that (4a) and (4b) hold. Then,

$$\sum_{qr} X_{qr}^2 = \text{tr} Y = k = \mathbf{1}^T X \mathbf{1} = \sum_{qr} X_{qr} = \sum_{qr} |X_{qr}|,$$

that is, $\|X\|_2 = \|X\|_1$. Combined with $0 \leq X_{qr} \leq 1$ this implies that $X_{qr} \in \{0, 1\}$. Since, in addition, $X \in \text{conv} \Pi_n^k$, we conclude that $X \in \Pi_n^k$. \square

To recap, so far we recasted the QAM problem (1) as the problem of maximizing a linear cost function (3a) subject to linear equalities and inequalities, (4a) and (4b), and the rank-one constraint (3b).

4.2. SDP relaxation

The next step of our construction is the relaxation of the rank-one constraint (3b). We replace $Y = [X][X]^T$ with $Y \succeq [X][X]^T$. That is, the constraint that $Y - [X][X]^T$ is a positive semidefinite (PSD) matrix. In turn, we devise the following optimization,

$$\max_{X,Y} \quad \text{tr}(WY) \quad (6a)$$

$$\text{s.t.} \quad Y \succeq [X][X]^T \quad (6b)$$

$$X \in \text{conv} \Pi_n^k \quad (6c)$$

$$\text{tr}Y = k \quad (6d)$$

Schur's complement (see *e.g.*, [BV04]) asserts that $Y \succeq [X][X]^T$ can be equivalently realized in semidefinite programming via the PSD constraint

$$\begin{bmatrix} Y & [X] \\ [X]^T & 1 \end{bmatrix} \succeq 0.$$

By construction (6) is a convex SDP relaxation of the QAM problem (1). In fact, it can be further improved (tightened), as we show in Section 4.4. This will be more straightforward once the relation to other relaxations is clear.

4.3. Relation to spectral and doubly-stochastic relaxations

Two relaxation approaches to the QAM problem (1) have become standard – spectral and doubly-stochastic relaxations (*e.g.*, [ABK14, LH05]). Next, we discuss the close relation of the SDP (6) with these relaxations. In particular, we show that, in a sense, it is equivalent to their “intersection” and therefore provably superior to both of them.

The doubly-stochastic approach relaxes the constraint (1b), that $X \in \Pi_n^k$, into its convex-hull $X \in \text{conv} \Pi_n^k$. Provided that $W \preceq 0$, the problem becomes a convex quadratic program [ABK14]. The spectral relaxation replaces $X \in \Pi_n^k$ with a necessary condition for k sub-permutations that $\|X\|_F^2 = \sum_{qr} X_{qr}^2 = k$. While still not convex, the resulting problem can be solved by simply taking the top eigenvector of the spectral decomposition of W [LH05]. As we prove next, both these standard relaxations are sub-optimal compared to the SDP relaxation (6), as illustrated in Figure 2.

Proposition 2 Both spectral and doubly-stochastic relaxations for the QAM problem (1) are sub-optimal compared to the SDP relaxation in (6). In fact, when formulated in the lifted variables, the feasible set of (6) is exactly the intersection of the feasible sets of these relaxations.

To prove this proposition, we first state two lemmas which provide a convex characterization of the spectral and doubly-stochastic relaxations. Their proofs can be found in Appendix A.

Consider the spectral relaxation (S), expressed in terms of the lift, and its relaxed counterpart (S') whereby $Y = [X][X]^T$ is replaced by $Y \succeq [X][X]^T$,

$$\begin{array}{ll} \max & \text{tr}WY \\ \text{(S)} \quad \text{s.t.} & \text{tr}Y = k \\ & Y = [X][X]^T \end{array} \quad \begin{array}{ll} \max & \text{tr}WY \\ \text{(S')} \quad \text{s.t.} & \text{tr}Y = k \\ & Y \succeq [X][X]^T \end{array}$$

Lemma 1 The relaxations (S) and (S') are equivalent.

Similarly, Consider the doubly-stochastic relaxation (DS), expressed in terms of the lift, and its relaxed counterpart (DS'),

$$\begin{array}{ll} \max & \text{tr}WY \\ \text{(DS)} \quad \text{s.t.} & X \in \text{conv} \Pi_n^k \\ & Y = [X][X]^T \end{array} \quad \begin{array}{ll} \max & \text{tr}WY \\ \text{(DS')} \quad \text{s.t.} & X \in \text{conv} \Pi_n^k \\ & Y \succeq [X][X]^T \end{array}$$

Lemma 2 The relaxations (DS) and (DS') are equivalent.

The proof of Proposition 2 now follows by noticing that the feasible set of (6) is the intersection of the feasible sets of (DS') and (S'). Figure 2 illustrates this intersection (in purple).

4.4. Tightening the relaxation

Thus far, we have shown that the convex SDP relaxation (6) is as good as both the spectral and doubly-stochastic relaxations. However, bridging the gap between these two relaxations in a convex optimization framework enables deriving an even better relaxation.

We propose adding linear equality and inequality constraints to the SDP relaxation in (6). Any constraints that are satisfied by the feasible points of the QAM (1) may be added; namely, constraints which hold for $Y = [X][X]^T$ with $X \in \Pi_n^k$. These constraints are guaranteed to improve (tighten) the relaxation provided that they are not containing the constraints of (6).

We devise the following convex SDP relaxation to the QAM problem (1), which we denote by QAM-SDP,

$$\max_Y \quad \text{tr}(WY) \quad (7a)$$

$$\text{s.t.} \quad Y \succeq [X][X]^T \quad (7b)$$

$$X \in \text{conv} \Pi_n^k \quad (7c)$$

$$\text{tr}Y = k \quad (7d)$$

$$Y \succeq 0 \quad (7e)$$

$$\sum_{qrst} Y_{qrst} = k^2 \quad (7f)$$

$$Y_{qrst} \leq \begin{cases} 0, & \text{if } q = s, r \neq t \\ 0, & \text{if } r = t, q \neq s \\ \min\{X_{qr}, X_{st}\}, & \text{otherwise} \end{cases} \quad (7g)$$

Here, the notation Y_{qrst} refers to double indexing of the entries of the matrix Y , corresponding to the lift $Y_{qrst} = X_{qr}X_{st}$. Note that QAM-SDP is a convex relaxation and can be plugged as-is into a semidefinite program (SDP) solver of choice. See additional implementation details in Section 7.

It is easy to verify that the additional linear equality and inequality constraints (7e)-(7g) hold for feasible points of

the QAM problem (1). Each constraint has an intuitive interpretation in terms of permutations. If $Y = [X][X]^T$ with $X \in \Pi_n^k$ then Y is a binary matrix with k^2 ones, which implies (7e) and (7f). To derive (7g), note that since X is a permutation its q 'th row has at most one non-zero entry, which entails that $Y_{qrst} = X_{qr}X_{qt} = 0$ for $r \neq t$; a similar argument for columns of X corresponds to the second case; otherwise, since X is binary, $Y_{qrst} = X_{qr}X_{st}$ implies that $Y_{qrst} \leq \min\{X_{qr}, X_{st}\}$.

Furthermore, the additional constraints (7e)-(7g) are not implied by the previous constraints (7b)-(7d). For instance, the latter do not imply that Y is non-negative: suppose $n = k$ and take X to be the $n \times n$ constant matrix whose entries equal $1/n$; let $u = [1, -1, 0, \dots, 0]^T$ and set

$$Y = [X][X]^T + \frac{n-1}{2}uu^T.$$

It follows that the pair (X, Y) satisfies (7b)-(7d). However, $Y_{12} = \frac{1}{n^2} - \frac{n-1}{2}$ is negative for $n > 1$, thus (7e) fails to hold. As such, these additional constraints are guaranteed to strictly tighten the QAM relaxation. This is summarized in the following proposition.

Proposition 3 The feasible set of the QAM-SDP relaxation (7) is *strictly* contained within the feasible set of relaxation (6). Therefore, it is *strictly* contained within the intersection of the feasible sets of the spectral (S') and doubly-stochastic (DS') relaxations.

Although this is not captured in full by Proposition 3, in practice the QAM-SDP relaxation (7) is far superior to the spectral and doubly-stochastic relaxations. As demonstrated in Section 6, the QAM-SDP relaxation often produces optimal or near-optimal results, which may be verified by checking that $Y = [X][X]^T$.

5. Consistent collection matching

Instead of matching just pairs of shapes, we may consider the problem of jointly matching a collection of shapes. As we show next, the QAM-SDP relaxation developed for matching pairs can be used to address this problem. Thus yielding an algorithm that combines, in a single optimization framework, a state-of-art pairwise matching algorithm and a global criterion enforcing compatibility between all pairwise matches.

Now, the input is a collection of m shapes $\mathcal{M}^1, \dots, \mathcal{M}^m$. The goal is to establish a good set of *consistent* partial matchings of order k between all pairs of shapes $(\mathcal{M}_i, \mathcal{M}_j)$, represented by sub-permutations $X^{ij} \in \Pi_n^k$.

A set of sub-permutations $\{X^{ij}\} \subset \Pi_n^k$ is *consistent* if it allows coloring a subset of k points in each shape such that X^{ij} relates points of the same color for any pair of shapes \mathcal{M}^i and \mathcal{M}^j , see for example the colored spheres in Figure 1. More formally, consistency allows well-defined labeling of the points of each shape in the set; inferring la-

bels from one shape to all others via X^{ij} is independent of the choice of a reference shape. The notion of consistency for permutations and functional maps was suggested and defined in [NBCW*11, HG13, HWG14]. It can be algebraically expressed as the transitivity relation $X^{ij}X^{j\ell} = X^{i\ell}$, for all i, j, ℓ .

5.1. Single-step consistent collection matching

Algorithms for matching a collection of shapes typically operate in two stages: first, matching many or all pairs of shapes in the collection, and second, combining all pairwise information by using a global principle, *e.g.*, by imposing map consistency. Examples include the works of [NBCW*11, SW12, KLM*12, HG13, HWG14]. However, altering correspondences to achieve global consistency may often lead to suboptimal matches.

We propose a single optimization step, in which pairwise matches are optimized in an inherent globally consistent manner. We formulate the problem of consistent collection matching (CCM) as the following global optimization problem,

$$\max_{\{X^{ij}\}} \sum_{i,j} [X^{ij}]^T W^{ij} [X^{ij}] \quad (8a)$$

$$\text{s.t. } X^{ij} \in \Pi_n^k \quad (8b)$$

$$\{X^{ij}\} \text{ are consistent} \quad (8c)$$

where $1 \leq i, j \leq m$. $W^{ij} \in \mathbb{R}^{n^2 \times n^2}$ defines a quadratic energy measuring the gain in matching \mathcal{M}_i and \mathcal{M}_j by a sub-permutation X^{ij} . It aims at establishing correspondences between all pairs of shapes, while respecting a global consistency requirement (8c).

5.2. Convex SDP relaxation.

Huang and Guibas [HG13] showed that a necessary and sufficient condition for the set $\{X^{ij}\}$ of *permutations* is that the block matrix containing all pairwise permutations,

$$\mathbf{X} = \begin{pmatrix} X^{11} & X^{12} & \dots & X^{1m} \\ X^{21} & X^{22} & \dots & X^{2m} \\ \vdots & \vdots & \ddots & \vdots \\ X^{m1} & X^{m2} & \dots & X^{mm} \end{pmatrix} \quad (9)$$

satisfies $\mathbf{X} \succeq 0$. We generalize their result to the case where \mathbf{X} comprises sub-permutations $\{X^{ij}\} \subset \Pi_n^k$. More specifically, in Appendix C we prove the following proposition.

Proposition 4 Assume that \mathbf{X} has sub-permutation blocks, $X^{ij} \in \Pi_n^k$. Then $\{X^{ij}\}$ are consistent if and only if $\mathbf{X} \succeq 0$.

Comparing Proposition 4 to the sample selection of Huang and Guibas in [HG13] reveals two differences: first, our formulation does not require fixing the subset of k points to be matched, for any of the shapes. Rather, this subset

is chosen automatically. Secondly, the linear constraints in their Proposition 2 seem unnecessary.

The CCM problem (8) is a generalization of the QAM (1) to the setting of collections. We therefore plug-in the QAM-SDP relaxation (7) into the CCM framework, as well as replace the consistency constraint (8c) with the PSD constraint $\mathbf{X} \succeq 0$. This leads to the CCM-SDP, the proposed relaxation of the CCM problem:

$$\max_{\mathbf{X}, \mathbf{Y}} \sum_{i,j} \text{tr}(W^{ij} Y^{ij}) \quad (10a)$$

$$\text{s.t.} \quad (X^{ij}, Y^{ij}) \in \mathcal{C}^k \quad \forall i < j \quad (10b)$$

$$X^{ii} \in \mathcal{D} \cap \text{conv} \Pi_n^k \quad \forall i \quad (10c)$$

$$\mathbf{X} \succeq 0 \quad (10d)$$

where \mathcal{C}^k is the set defined by the convex SDP constraints (7b)-(7g), with k indicating the sub-permutation order. X^{ii} is the mapping of the shape \mathcal{M}^i onto itself, and is therefore a diagonal matrix (denoted by \mathcal{D}) that is also in $\text{conv} \Pi_n^k$. Intuitively, the non-zero values on the diagonal of X^{ii} identify the subset of k points in shape \mathcal{M}^i that are matched. Lastly, $\mathbf{Y} = \{Y^{ij}\}$ is the set of all lifted sub-permutations between each pair of shapes \mathcal{M}^i and \mathcal{M}^j .

As demonstrated in Section 6, the CCM-SDP relaxation often produces optimal or near-optimal results. Optimality can be verified by checking that $Y^{ij} = [X^{ij}][X^{ij}]^T$ for all i, j . In this case, a consistent set of sub-permutations can be readily extracted from \mathbf{Y} or \mathbf{X} . Appendix D details a projection scheme onto the space of consistent sub-permutations, addressing cases in which the optimization yields sub-optimal results.

6. Experimental results

We have experimented with our convex QAM-SDP and CCM-SDP relaxations for matching sets of points within pairs and collections of shapes originating from synthetic, computer graphics and computational anatomy datasets. In all cases, the suggested convex relaxations proved to be extremely tight, in the sense that very often they returned the certified globally optimal solution to these challenging optimization problems.


6.1. Synthetic evaluation of QAM-SDP for pairwise matching

We have tested the QAM-SDP relaxation independently against three relaxations of the QAM problem: the spectral and DS relaxations described above, and a straightforward SDP relaxation, that we name D-SDP. Following a similar derivation to that of [KSSC03], this relaxation is obtained by directly computing the dual of the binary QAM problem, which also leads to an SDP.

Each relaxation provides an upper-bound p^+ to the optimal value p^* of the QAM problem. We project the relaxed result onto Π_n^k using linear-programming over the set of doubly stochastic matrices. (For the D-SDP method we use the Lagrangian to extract a corresponding primal solution and then project.) This provides a feasible instance X of the QAM and therefore a lower bound p^- to its optimal value. Thus, each method provides bounds on the optimal value $p^- \leq p^* \leq p^+$.

We measure the quality of each relaxation method by computing its *relative optimality gap*, defined by $\Delta = (p^+ - p^-)/p^- \geq 0$. This indicates the optimality of a solution X , and in particular, vanishes for the global optimum of the QAM. We further compare the objective value p^- of a projected solution of each of the methods to QAM-SDP via the *objective ratio* $\Gamma = p_{\text{QAM-SDP}}^-/p^-$. Here, $\Gamma > 1$, represented as percentage in the figures, implies that QAM-SDP has achieved higher (*i.e.*, better) objective value.

Evaluation. We evaluated the matching performance of each of the methods over random pairs of shapes, generated as follows. We sampled $n = 10$ points uniformly at random on a sphere. We deformed the sphere with an increasing level of noise to create a set of 20 shapes $\mathcal{M}^1, \mathcal{M}^1, \mathcal{M}^2, \dots, \mathcal{M}^{20}$, as the inset illustrates. We then compared \mathcal{M}^1 to each of the noise perturbed shapes. We repeated this experiment 100 times, for an overall of 2000 matching experiments. For this experiment we took $k = n$, corresponding to the case of optimization over (full) permutations.



We have considered two instantiations of the QAM problem (1), that is, choices of W :

1. Graph matching functional [ABK14] – relaxation prescribed by

$$[X]^T W [X] = -\|XD^1 - D^2X\|_F^2,$$

where D^1, D^2 are the Euclidean distance matrices of each of the two compared shapes, that is, $(D^1)_{qr} = \|\mathbf{p}_q^1 - \mathbf{p}_r^1\|_2$, and similarly for D^2 . The expression for the corresponding W is provided for completeness in Appendix B.

2. Isometric matching functional – use Eq. (2) to define the QAM objective. We choose σ to be proportional to the standard deviation of the absolute values of all the pairwise differences between entries of D^1 and D^2 .

Note that the first case defines a concave quadratic functional, as $W \preceq 0$, and therefore DS is suitable form of relaxation in this case. For the isometric matching functional, however, $W \not\preceq 0$ thus the DS relaxation cannot be directly employed. On the other hand, using the spectral relaxation is natural in this case. Since W is entry-wise non-negative

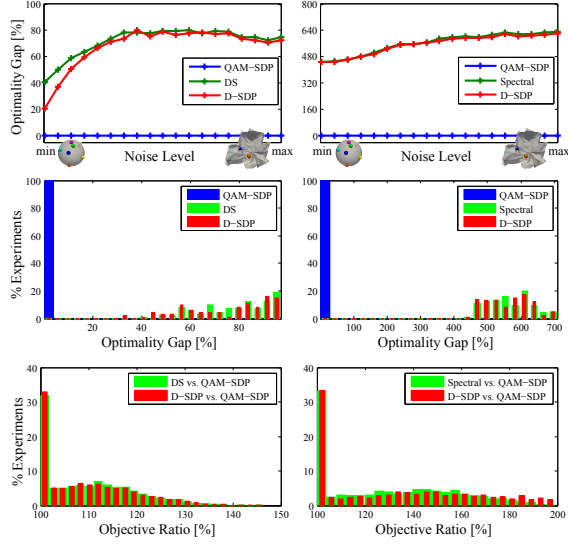


Figure 3: Evaluation of QAM relaxations for pairwise matching. Left – graph matching functional; Right – isometric matching functional. First row – average optimality gaps as a function of perturbation level. Second and third rows – distribution of optimality gaps and objective ratios (for medium noise level).

the Perron-Frobenius theorem [HJ90] guarantees that the entries of the top eigenvector are all within the interval $[0, 1]$. In both cases we also applied the D-SDP.

Figure 3 summarizes the results of these experiments where the left column shows the first instantiation and the right column the second one. The first row shows the average optimality gaps as a function of deformation level (with each data point representing 100 experiments).

The second and third rows show statistics of 100 experiments performed with medium noise level. The second row shows the histogram of relative optimality gaps of each of the methods and the third shows the histogram of the objective ratios of the D-SDP, spectral and doubly-stochastic relaxations compared to QAM-SDP. (Values over 100% indicate that QAM-SDP achieved higher objective in its solution.) These experiments demonstrate that the quality of the different relaxations deteriorates as the perturbation level increases, while QAM-SDP produces optimal results (near zero optimality gap). This behaviour persists even at very high levels of noise. The histograms of objective ratios indicate that the functional value of the different relaxation methods is equal or higher than the values produced by the QAM-SDP relaxation.

As an additional experiment, we evaluated the performance of QAM-SDP with a random objective matrix W , at different problem dimensions $n = 2, \dots, 15$. W was generated by drawing its entries from normal distribution followed




Figure 4: Evaluation of QAM relaxations for random objectives. Left – average optimality gaps vs. problem size (n). Right – distribution of optimality gaps (for $n = 10$); the D-SDP optimality gaps are concentrated above 100%, illustrated by the blowup.

by rank projection via SVD; the rank was drawn uniformly at random between zero to n . We repeated this experiment 100 times for each problem dimension n , for an overall of 1400 matching experiments. We evaluated the performance of QAM-SDP and D-SDP. Figure 4 (left) summarizes the results of this experiment, showing the average optimality gaps Δ_{D-SDP} and $\Delta_{QAM-SDP}$ as a function of problem size n . In this scenario, as the problem size increases, QAM-SDP does not achieve zero optimality gap. Nonetheless, it is significantly superior to D-SDP. This is also demonstrated in Figure 4 (right), which shows the distribution of optimality gaps for $n = 10$.

6.2. Matching pairs of shapes

Evaluating the performance of QAM-SDP in a non-synthetic scenario is not straightforward. Existing evaluation protocols for shape matching algorithms assume a complete shape matching pipeline, and are not suitable for independently evaluating our proposed relaxation for the QAM problem as formulated in Section 4. We therefore used simple pre- and post-processing components to obtain a basic complete shape matching pipeline.


We have experimented with intra-class pairs of shapes taken from the SHREC dataset [GBP07]. For each shape we automatically generated a sparse set of candidate points. We followed [KLCF10] by taking the local critical points of the approximate average geodesic distance (AGD) function, which typically contain a subset of semantically meaningful points. We completed a full set of n candidate points by farthest point sampling. Both n and k were manually chosen. We further associate with \mathcal{M}_i an approximate geodesic distance d_i , approximated by computing Dijkstra's shortest path length, normalized by the approximate geodesic diameter of the shape. We used the isometric matching functional, Eq. (2), to set W .

Correspondences generated by our approaches for various pairs of shapes are presented in Figure 5. Note that the QAM-SDP requires no initialization nor employs additional information regarding the shapes except the distances between the feature points. For the majority of




Figure 5: Quadratic Assignment Matching. Coarse correspondences found using the QAM-SDP relaxation applied to pairs of shapes of SHREC 2007 database. A simple procedure was used to find candidate feature points. Correspondences computed with QAM-SDP are visualized by colored spheres. Surface colors visualize dense maps induced by these sparse correspondences.

experiments, QAM-SDP achieved an optimal result, as summarized in the histogram of optimality gaps shown in the inset. In particular, our pair matching experiments comprised over 7000 QAM-SDP optimizations, of which over 90% achieved the *global optimum* for the corresponding QAM problem.



For completeness, we further followed the evaluation protocol of [KLF11], indicating the quality of *dense* vertex-vertex correspondences. It computes the percentage of correspondences whose deviations from ground-truth correspondences fall within a varying threshold. To this end, we used post-processing to induce vertex-vertex correspondences between pairs of models: pairs were bijectively mapped using Tutte's embedding to the unit disc, with the additional requirement that corresponding points are mapped to each other; then, naive vertex-vertex correspondences were deduced by searching nearest-neighbors in the plane. The surface colors in Figure 5 visualize the obtained maps, indeed interpolating the colors of the correspondences. The inset shows comparable results to those reported by [KLF11] for the same categories and pairs of shapes.



6.3. Synthetic evaluation of CCM-SDP for collection matching

Evaluation of randomly generated collections. We evaluated the performance of the one-stage CCM-SDP algorithm compared to two-stage algorithms, that first compute the pairwise matchings $X^{ij} : \mathcal{M}^i \leftrightarrow \mathcal{M}^j$ and then project \mathbf{X} onto the space of consistent sub-permutations. We compared to the projection method of Huang and Guibas [HG13] (HG), as well as to the heuristic projection (Proj) described in Appendix D. We feed both methods with X^{ij} computed with our QAM-SDP relaxation, which precisely corresponds to the pairwise matching component in the CCM-SDP algorithm.

For each of three overall noise levels, we generated 500 sequences of 10 increasingly deforming shapes $\mathcal{M}^1, \dots, \mathcal{M}^{10}$, as described in Section 6.1. We matched $n = k = 8$ points, and used the isometric matching functional, Eq. (2), to set W^{ij} for every pair of shapes \mathcal{M}^i and \mathcal{M}^j .

Figure 6 summarizes the results of this experiment. The top row shows the distributions of relative optimality gaps and objective values, obtained with HG and Proj relative to CCM-SDP. The latter presents far lower optimality gaps compared to other methods. In particular, in more than 70% of the experiments CCM-SDP achieved globally optimal results (zero gap). Moreover, in the majority of experiments the single-stage CCM-SDP achieved higher objective value (p^*) compared to the two-stage methods.

To further demonstrate the advantage of the single-stage CCM-SDP approach we evaluated the quality of the result-

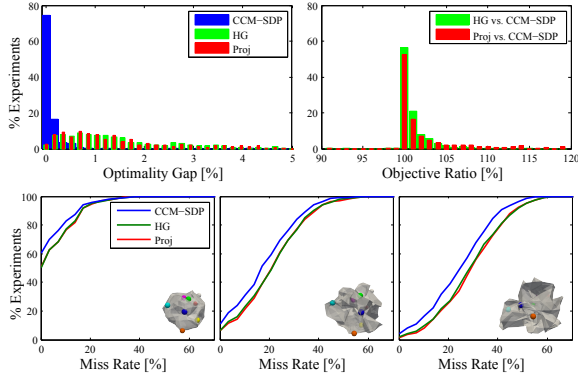


Figure 6: Evaluation of CCM relaxations for matching collections. Top – distribution of optimality gaps and objective ratios (for high noise level). Bottom – miss-rate statistics, for three increasing levels of deformation. HG and Proj are initialized with pairwise correspondences computed with QAM-SDP.

ing correspondences. For each collection of 10 shapes we computed the ratio of correct matchings, over all pairs of shapes \mathcal{M}^i and \mathcal{M}^j , with respect to the ground-truth labeling of the points. The bottom of Figure 6 shows the percentage of experiments whose miss-rate (percentage of points incorrectly matched) falls within a varying threshold. Each graph summarizes 500 experiments conducted on sequences of increasingly deformed shapes, at three noise levels. The insets illustrate the highest noise level of the data corresponding to each graph. In these experiments, the CCM-SDP consistently produces lower miss-rates. The single-step CCM-SDP relaxation has a significant advantage in scenarios involving large shape variability. This advantage lessens at low levels of deformation, as HG and Proj are already initialized with very good initial pairwise matchings provided by the QAM-SDP.

Synthetic evaluation of shape collections. We further evaluated the performance of the single-step CCM-SDP compared to two-step approaches on 2 classes of models from the 3D Warehouse (chairs and planes), for which a sparse set of ground-truth correspondences is available [KLM*13].

For each class, we randomly select 10 distinct shapes $\mathcal{M}^1, \dots, \mathcal{M}^{10}$. We use the feature points provided with the datasets, 10 points for chairs and 7 for the planes and bikes. We chose n accordingly. In all cases we took $n = k$. As before, we used the isometric matching functional, Eq. (2), to set W^{ij} for every pair of shapes \mathcal{M}^i and \mathcal{M}^j . We repeated each experiment 100 times.

Figure 7 summarizes the results of these experiments. The miss-rate statistics, distribution of optimality gaps and objective ratios are presented for each class, as previously de-

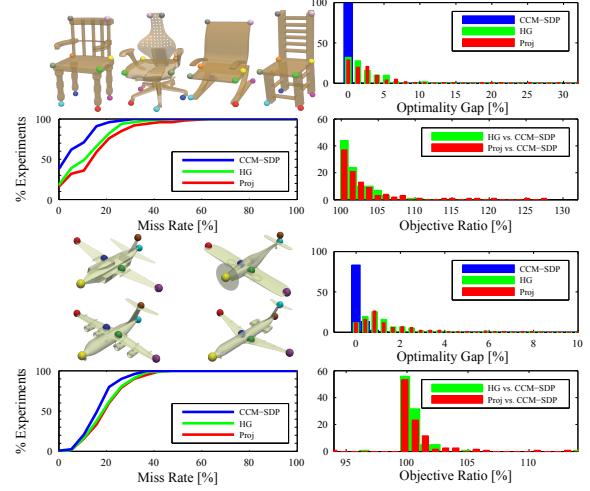


Figure 7: Evaluation of CCM relaxations for matching collections of shapes. For each set (chairs, planes) – miss-rate statistics, distribution of optimality gaps and objective ratios. HG and Proj are initialized with pairwise correspondences computed with QAM-SDP. Representative models from set show the correspondences established by CCM-SDP.

scribed. In addition, for each class a selection of models from a typical random collection is presented, overlaid with the correspondences established by CCM-SDP.

These experiments suggest that establishing pairwise correspondences in an inherently consistent manner is preferable to two-step alternatives. Intuitively, enforcing consistency *during* the optimization of correspondences between pairs has an “error-correction” effect: stable correspondences (e.g., for similar shapes) propagate throughout the collection, eventually disambiguating pairs of weaker relations. Therefore, improving robustness to large variations within the collection.

6.4. Matching collections of shapes

Matching collections of SHREC models. We employed the CCM-SDP framework to establish consistent correspondences over a few categories of similar models taken from the SHREC 2007 dataset [GBP07], each containing about 20 models. We followed the procedure described in Section 6.2 for generating feature points and setting W^{ij} for every pair of shapes. Figures 1 and 8 present sparse consistent correspondences automatically established using our approach. Note that our method succeeds in computing high quality consistent correspondences over each of the collections, inspite the presence of significant shape variability.

Matching collections of anatomical models. We used the dataset of [BLSC*11], comprising 94 surfaces representing




Figure 8: Consistent Collection Matching. Consistent sparse correspondences found using the CCM-SDP relaxation applied to two classes of SHREC 2007 database are visualized by colored spheres. Surface colors visualize induced dense maps.

high-resolution digitizations of teeth. This is a highly challenging dataset that consists of non-isometric models with large shape and feature variations. It is of particular interest in this type of anatomical data to identify consistent landmarks across the collection. This problem directly fits into our CCM framework that is guaranteed to return consistently labeled points.

For this application, we used Euclidean distances for the construction of the pairwise functionals W^{ij} . To choose candidate points, we flattened each model by minimizing the LSCM functional [LPRM02]; we then chose the local minima of the conformal factor, indicating points of maximal local shrinkage. As suggested in [BLSC*11], these points demonstrate good correspondence with morphologically salient points. We completed to a full set of n candidate points by farthest point sampling.

We employed the CCM-SDP relaxation to collections of 20 models selected from the entire dataset. Figure 9 shows the consistent correspondences found by our approach. We further used the results to extrapolate dense correspondence maps. To this end, we registered the flattened versions of pairs of models using thin plate splines, aiming to match the correspondences found by CCM-SDP; then, vertex-vertex correspondences are deduced by nearest-neighbor search. The surface colors in Figure 9 visualize the obtained maps. Lastly, in the inset we compare our results against ground truth correspondences (human-marked landmarks, provided with the dataset). Note that more than 80% of the





Figure 9: Consistent sparse correspondences computed by CCM-SDP on a collection of 20 surfaces from a dataset of teeth models [BLSC*11]. Dense maps, visualized by surface color, are inferred from these correspondences.

ground-truth correspondences were found with distance error smaller than 10% of the diameter.

6.5. Collection metric learning

The CCM framework naturally establishes a metric or an affinity measure on the shape space $\{\mathcal{M}^i\}$. In this application we used the CCM-SDP algorithm to automatically produce an affinity measure on the collection of teeth surfaces studied in the previous section, and employed mani-






Figure 10: 3D embedding of the collection obtained by applying diffusion maps to pairwise affinities provided by our CCM-SDP method. The Euclidean coordinates correlate with biological classification.

fold learning technique in order to embed it into a lower dimensional Euclidean space. The Euclidean coordinates correlated nearly perfectly with ground truth biological classification.

Specifically, the functional of the CCM problem, (8a), is of the form of a sum $\sum_{ij} a_{ij}$ where

$$a_{ij} = [X^{ij}]^T W^{ij} [X^{ij}]$$

quantifies the affinity between the models \mathcal{M}_i and \mathcal{M}_j . Intuitively, a higher a_{ij} value indicates that better correspondences were found (with respect to the problem model W^{ij}).

We computed the affinities associated with our CCM-SDP solution for the teeth models shown in Figure 9. Then, we used diffusion maps [NLCK05] to find a 3D embedding of the collection, realizing the diffusion distances. Figure 10 illustrates this embedding; the 3D embedded coordinates of each of the models are represented by both their spatial locations as well as their RGB colors.

We further validate the results by comparing to ground truth classification provided with the data [BLSC*11]. Figure 11 shows the same models, with the same coloring. The underlying color textures indicate the ground truth classification of the collection according to genus.

7. Implementation details and limitations

We implemented the QAM-SDP and CCM-SDP algorithms in MATLAB, using YALMIP for the modeling of semidefinite programs [L04].

The main limitation of the proposed framework is its computational scalability. The number of variables in the QAM-SDP and CCM-SDP relaxations are $o(n^4)$ and $o(m^2n^4)$, respectively; thus the problem size scales quadratically with the number of models m and quartically with the dimension of sub-permutations n . Nevertheless, it provides a polynomial-time approach for approximating an NP-hard problem, which is considered to be large scale for $n \geq 16$ and generally intractable for $n \geq 30$ [ZKRW98, LdABN*07].





Figure 11: Comparison to ground truth classification of genus. Models are colored by their 3D embedding. The ground truth classes are indicated by the underlying texture.

As of today, interior point methods are the standard approach for solving SDP with high accuracy. Unfortunately, these are second order methods which tend to scale unfavorably with problem size. We therefore used MOSEK [AA99], an interior point solver, for computing the certified global optimum of relaxations with $n \leq 15$ within acceptable times. We have further experimented with a first order ADMM solver, SCS [OCBP13], for larger problems. Although first order methods scale well with problem size, they are known to only achieve moderate accuracy within a reasonable number of iterations; we terminated this solver after 5000 iterations. Figure 12 presents QAM-SDP results obtained using ADMM optimization for three pairs of shapes taken from the SHREC dataset. The inset shows typical run-times for solving pairwise problems of varying dimension, optimized using MOSEK or SCS; timings were measured on a 2.50GHz Intel Xeon.



8. Concluding remarks

In this paper, we have considered the combinatorial problem of matching given sets of points, modeled as a Quadratic Assignment Matching (QAM) problem. We have presented a novel convex relaxation for this NP-hard problem that builds upon a rank-one reformulation of the problem in a higher dimension. The proposed formulation is shown to be a hybrid of the standard spectral and doubly-stochastic relaxations. In particular, we prove that it is in fact tighter than both.

Our approach is further generalized to the problem of Consistent Collection Matching (CCM). We introduce a convex relaxation which enables establishing consistent correspondences over a collection of shapes, in a single optimiza-




Figure 12: Quadratic Assignment Matching using ADMM for a larger number of points ($20 \leq n \leq 30$).

tion. This is in contrast to alternative approaches, which first estimate pairwise correspondences and then enforce global consistency in a second step.

Experiments show that these relaxations are extremely tight, often producing guaranteed optimal solution for the QAM and CCM problem, far better than existing approaches. Thus, providing a well performing polynomial-time approximation for a notoriously hard problem. This, however, comes at the cost of a solving an SDP in a higher dimension, which limits the scalability of the approach. The properties of the proposed approach, as well as the unified convex perspective it provides, offer a contribution that goes beyond the practical limitations dictated by current evolving optimization technology.

Acknowledgements This work was supported in part by the European Research Council (ERC Starting Grant, grant No. 307754 "SurfComp"), the Israel Science Foundation (grant No. 1284/12 and 1265/14) and the I-CORE program of the Israel PBC and ISF (Grant No. 4/11). The authors would like to thank Amit Singer for useful discussions and the anonymous reviewers for their helpful comments and suggestions.

References

- [AA99] ANDERSEN E. D., ANDERSEN K. D.: *The MOSEK interior point optimization for linear programming: an implementation of the homogeneous algorithm*. Kluwer Academic Publishers, 1999, pp. 197–232. [11](#)
- [ABK14] AFLALO Y., BRONSTEIN A., KIMMEL R.: Graph matching: relax or not? *arXiv preprint arXiv:1401.7623* (2014). [2, 3, 4, 6](#)
- [BBK06] BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R.: Generalized multidimensional scaling: A framework for isometry-invariant partial surface matching. *Proceedings of the National Academy of Sciences of the United States of America* 103, 5 (2006), 1168–1172. [1, 2, 3](#)
- [BBK*10] BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R., MAHMOUDI M., SAPIRO G.: A gromov-hausdorff framework with diffusion geometry for topologically-robust non-rigid shape matching. *International Journal of Computer Vision* 89, 2–3 (2010), 266–286. [2](#)
- [BBM05] BERG A. C., BERG T. L., 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* (2005), vol. 1, IEEE, pp. 26–33. [1, 2, 3](#)
- [BLSC*11] BOYER D. M., LIPMAN Y., ST. CLAIR E., PUENTE J., PATEL B. A., FUNKHOUSER T., JERNVALL J., DAUBECHIES I.: Algorithms to automatically quantify the geometric similarity of anatomical surfaces. *Proceedings of the National Academy of Sciences* 108, 45 (2011), 18221–18226. [9, 10, 11](#)
- [BV04] BOYD S., VANDENBERGHE L.: *Convex Optimization*. Cambridge University Press, 2004. [4, 13](#)
- [DBKP11] DUCHENNE O., BACH F., KWEON I.-S., PONCE J.: A tensor-based algorithm for high-order graph matching. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* 33, 12 (2011), 2383–2395. [2](#)
- [DM58] DULMAGE A., MENDELSON N.: The convex hull of sub-permutation matrices. In *Proc. Amer. Math. Soc* (1958), vol. 9, pp. 253–254. [3](#)
- [FHJB13] FENG W., HUANG J., JU T., BAO H.: Feature correspondences using morse small complex. *The Visual Computer* 29, 1 (2013), 53–67. [2, 3](#)
- [FS06] FUNKHOUSER T., SHILANE P.: Partial matching of 3d shapes with priority-driven search. In *Proceedings of the Fourth Eurographics Symposium on Geometry Processing* (Aire-la-Ville, Switzerland, Switzerland, 2006), SGP '06, Eurographics Association, pp. 131–142. [1, 2](#)
- [GBP07] GIORGI D., BIASOTTI S., PARABOSCHI L.: Shape retrieval contest 2007: Watertight models track, 2007. [2, 7, 9](#)
- [HG13] HUANG Q.-X., GUIBAS L.: Consistent shape maps via semidefinite programming. *Computer Graphics Forum* 32, 5 (2013), 177–186. [2, 5, 8](#)
- [HJ90] HORN R. A., JOHNSON C. R.: *Matrix Analysis*. Cambridge University Press, 1990. [7](#)
- [HWG14] HUANG Q., WANG F., GUIBAS L.: Functional map networks for analyzing and exploring large shape collections. *ACM Trans. Graph.* 33, 4 (July 2014), 36:1–36:11. [2, 5](#)
- [KLCF10] KIM V. G., LIPMAN Y., CHEN X., FUNKHOUSER T.: Möbius transformations for global intrinsic symmetry analysis. In *Computer Graphics Forum* (2010), vol. 29, Wiley Online Library, pp. 1689–1700. [7](#)
- [KLF11] KIM V. G., LIPMAN Y., FUNKHOUSER T.: Blended intrinsic maps. In *ACM Transactions on Graphics (TOG)* (2011), vol. 30, ACM, p. 79. [8](#)
- [KLM*12] KIM V. G., LI W., MITRA N. J., DiVERDI S., FUNKHOUSER T.: Exploring Collections of 3D Models using Fuzzy Correspondences. *Transactions on Graphics (Proc. of SIGGRAPH)* 31, 4 (2012). [2, 5](#)
- [KLM*13] KIM V. G., LI W., MITRA N. J., CHAUDHURI S., DiVERDI S., FUNKHOUSER T.: Learning Part-based Templates from Large Collections of 3D Shapes. *Transactions on Graphics (Proc. of SIGGRAPH)* 32, 4 (2013). [9](#)
- [KSSC03] KEUCHEL J., SCHNORR C., SCHELLEWALD C., CREMERS D.: Binary partitioning, perceptual grouping, and restoration with semidefinite programming. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* 25, 11 (2003), 1364–1379. [6](#)
- [LÖ4] LÖFBERG J.: Yalmip : A toolbox for modeling and optimization in MATLAB. In *Proceedings of the CACSD Conference* (Taipei, Taiwan, 2004). [11](#)

- [LdABN*07] LOIOLA E. M., DE ABREU N. M. M., BOAVENTURA-NETTO P. O., HAHN P., QUERIDO T.: A survey for the quadratic assignment problem. *European Journal of Operational Research* 176, 2 (2007), 657 – 690. 2, 11
- [LH05] LEORDEANU M., HEBERT M.: A spectral technique for correspondence problems using pairwise constraints. In *Computer Vision, 2005. ICCV 2005. Tenth IEEE International Conference on* (2005), vol. 2, IEEE, pp. 1482–1489. 2, 3, 4
- [LPRM02] LÉVY B., PETITJEAN S., RAY N., MAILLOT J.: Least squares conformal maps for automatic texture atlas generation. *ACM Trans. Graph.* 21, 3 (July 2002), 362–371. 10
- [Mem07] MEMOLI F.: On the use of gromov-hausdorff distances for shape comparison. In *Eurographics symposium on point-based graphics* (2007), The Eurographics Association, pp. 81–90. 2
- [NBCW*11] NGUYEN A., BEN-CHEN M., WELNICKA K., YE Y., GUIBAS L.: An optimization approach to improving collections of shape maps. *Computer Graphics Forum* 30, 5 (2011), 1481–1491. 2, 5
- [NLCK05] NADLER B., LAFON S., COIFMAN R. R., KEVREKIDIS I. G.: Diffusion maps, spectral clustering and eigenfunctions of fokker-planck operators. *arXiv preprint math/0506090* (2005). 11
- [OCPB13] O'DONOGHUE B., CHU E., PARIKH N., BOYD S.: Operator splitting for conic optimization via homogeneous self-dual embedding. *arXiv preprint arXiv:1312.3039* (2013). 11
- [SLZ*13] SHAO T., LI W., ZHOU K., XU W., GUO B., MITRA N. J.: Interpreting concept sketches. *ACM Transactions on Graphics (TOG)* 32, 4 (2013), 56. 2
- [SW12] SINGER A., WU H.-T.: Vector diffusion maps and the connection laplacian. *Communications on Pure and Applied Mathematics* 65, 8 (2012), 1067–1144. 2, 5
- [vKZHC011] VAN KAICK O., ZHANG H., HAMARNEH G., COHEN-OR D.: A survey on shape correspondence. *Computer Graphics Forum* 30, 6 (2011), 1681–1707. 2
- [ZKRW98] ZHAO Q., KARISCH S. E., RENDL F., WOLKOWICZ H.: Semidefinite programming relaxations for the quadratic assignment problem. *Journal of Combinatorial Optimization* 2, 1 (1998), 71–109. 2, 11

A. Convex representation of standard relaxations

Proof of Lemma 1 To prove the equivalence of (S') and (S) we show that the maximizer of (S') is smaller or equal to the maximizer of (S) . Denote by μ_1 the top eigenvalue of W . In turn, note that $k\mu_1$ is exactly the optimal objective value of (S) . Let u_1, u_2, \dots, u_{n^2} and $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_{n^2}$ be the eigenvectors and eigenvalues of Y . Then,

$$\text{tr} WY = \sum_{\ell} \lambda_{\ell} u_{\ell}^T W u_{\ell} \leq \mu_1 \sum_{\ell} \lambda_{\ell} = k\mu_1,$$

where we used the fact that $\max_{u: \|u\|_2=1} u^T W u = \mu_1$, and that $\sum_{\ell} \lambda_{\ell} = \text{tr} Y = k$. \square

Proof of Lemma 2 Since (DS') is a relaxation of (DS) it is enough to show that the maximizer of (DS') is smaller or equal to the maximizer of (DS) . Indeed, for $W \preceq 0$, where the (DS) relaxation is well posed, the constraint $Y - [X][X]^T \succeq 0$ implies that $\text{tr} WY \leq \text{tr} W[X][X]^T$. \square

B. Graph matching functional

Recall that $[AXB] = (B^T \otimes A)[X]$, where \otimes is the Kronecker product of matrices. Therefore, for a permutation matrix X we have

$$[X]^T W [X] = -\|XD^1 - D^2 X\|_F^2,$$

with

$$W = -\left(D^1{}^T \otimes I - I \otimes D^2\right)^T \left(D^1{}^T \otimes I - I \otimes D^2\right).$$

C. Consistency constraint

We prove Proposition 4. We make use of a Schur complement (with a singular block) formulated below for completeness (see, e.g., [BV04], p. 651).

The direction ' \Rightarrow ' follows the fact that consistency implies that for all i, j , $X^{i1}X^{1j} = X^{ij}$ and therefore, by construction, $\mathbf{X} = \begin{bmatrix} X^{11} & X^{12} & \dots & X^{1m} \end{bmatrix}^T \begin{bmatrix} X^{11} & X^{12} & \dots & X^{1m} \end{bmatrix} \succeq 0$.

To prove the other direction ' \Leftarrow ' assume that $\mathbf{X} \succeq 0$. First, we observe that its diagonal blocks $X^{\ell\ell}$ are diagonal; otherwise, since $X^{\ell\ell} = (X^{\ell\ell})^T \in \Pi_n^k$, it has a principal minor $\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$ in contradiction to $\mathbf{X} \succeq 0$.

Next, we need to show that $X^{ij} = X^{i\ell}X^{\ell j}$ for all i, j and ℓ . Without loss of generality, it suffices to show that $X^{ij} = X^{i1}X^{1j}$. To this end, we Partition \mathbf{X} as follows

$$\mathbf{X} = \begin{bmatrix} X^{11} & \mathbf{X}^{1,2:m} \\ (\mathbf{X}^{1,2:m})^T & \mathbf{X}^{2:m,2:m} \end{bmatrix} \succeq 0, \quad (11)$$

with the notation

$$\mathbf{X}^{q:r,s:t} = \begin{bmatrix} X^{q,s} & \dots & X^{q,t} \\ \vdots & \ddots & \vdots \\ X^{r,s} & \dots & X^{r,t} \end{bmatrix}.$$

Schur's complement then implies that

$$\mathbf{X}^{2:m,2:m} - (\mathbf{X}^{1,2:m})^T (X^{11})^{\dagger} \mathbf{X}^{1,2:m} \succeq 0 \quad (12)$$

$$\mathbf{X}^{1,2:m} \in \mathcal{R}(X^{11}) \quad (13)$$

where $(X^{11})^{\dagger}$ is the pseudo-inverse of X^{11} , and $\mathcal{R}(X^{11})$ denotes the column space of the matrix X^{11} .

Since X^{11} is diagonal with $\{0, 1\}$ entries we have $(X^{11})^{\dagger} = X^{11}$. Eq. (13) entails that $X^{11}X^{1\ell} = X^{1\ell}$, for $\ell = 2, \dots, m$. Combining these we obtain

$$(X^{1\ell})^T (X^{11})^{\dagger} X^{1\ell} = (X^{1\ell})^T X^{11} X^{1\ell} = (X^{1\ell})^T X^{1\ell} = Q_{\ell},$$

for every $\ell = 2, \dots, m$, where Q_{ℓ} is a diagonal matrix with exactly k ones on its diagonal. For the ℓ 'th diagonal block of (12) we have

$$X^{\ell\ell} - Q_{\ell} \succeq 0. \quad (14)$$

Since $X^{\ell\ell}$ is diagonal as well as in Π_n^k , it also has exactly k ones on its diagonal. The only case Eq. (14) can hold is if $X^{\ell\ell} = Q_\ell$. Thus, the left-hand-side in (12) is a PSD matrix whose diagonal blocks are all zeros, which implies it must be the zero matrix. In turn, (12) and (13) reduce to

$$\mathbf{X}^{2:m,2:m} = (\mathbf{X}^{1,2:m})^T \mathbf{X}^{1,2:m},$$

whose blocks entail that $X^{ij} = X^{i1} X^{1j}$ for $i, j = 2, \dots, m$. The case where $i = 1$ or $j = 1$ can be easily verified to hold using (13), which completes the proof. \square

Lemma 3 (Schur complement with a singular block) Let

$X = \begin{bmatrix} A & B \\ B^T & C \end{bmatrix}$ be a symmetric matrix. Then,

$$X \succeq 0 \iff A \succeq 0, B \in \mathcal{R}(A), C - B^T A^\dagger B \succeq 0.$$

D. Projection on the set of consistent sub-permutations.

When the optimization of (10) returns near- or sub-optimal solution we project the resulting \mathbf{Y} onto the space of consistent sub-permutations, according to the following procedure:

1. Factorize Y^{ij} – use eigen-decomposition to compute the best rank-one factorization of Y^{ij} ; that is, find $[X^{ij}]$ that best approximates $Y^{ij} \approx [X^{ij}] [X^{ij}]^T$. The sign ambiguity in $[X^{ij}]$ is resolved by majority voting. The diagonal blocks X^{ii} are taken directly from the optimization output. Combine the resulting X^{ij} to form a new matrix $\hat{\mathbf{X}}$.
2. Factorize $\hat{\mathbf{X}}$ – use eigen-decomposition to compute the best rank- k factorization of $\hat{\mathbf{X}}$; that is, find $U \in \mathbb{R}^{mn \times k}$ such that $\hat{\mathbf{X}} \approx UU^T$. To resolve the ambiguity in this factorization, we use the fact that we know that each of the $n \times k$ blocks of U should reside in $\Pi_{n \times k}^k$, the set of rectangular $n \times k$ sub-permutation matrices of order k . Let \tilde{U}_1 be the $k \times k$ matrix obtained by taking the k rows of largest norm out of the n first rows of U . Then set $V = U\tilde{U}_1^{-1}$.
3. Project onto sub-permutations – form a matrix \tilde{V} by projecting each $n \times k$ block of V onto $\Pi_{n \times k}^k$. The projection is done using linear programming. As numerical errors sometimes prevent perfect integer solutions, few steps of branch and bound are guaranteed to provide the globally optimal integer solution at almost no extra computational effort. The final consistent sub-permutations are the $n \times n$ blocks of $\mathbf{X} = \tilde{V}\tilde{V}^T$.