Point Registration via Efficient Convex Relaxation

Haggai Maron Nadav Dym Itay Kezurer Shahar Kovalsky Yaron Lipman
Weizmann Institute of Science

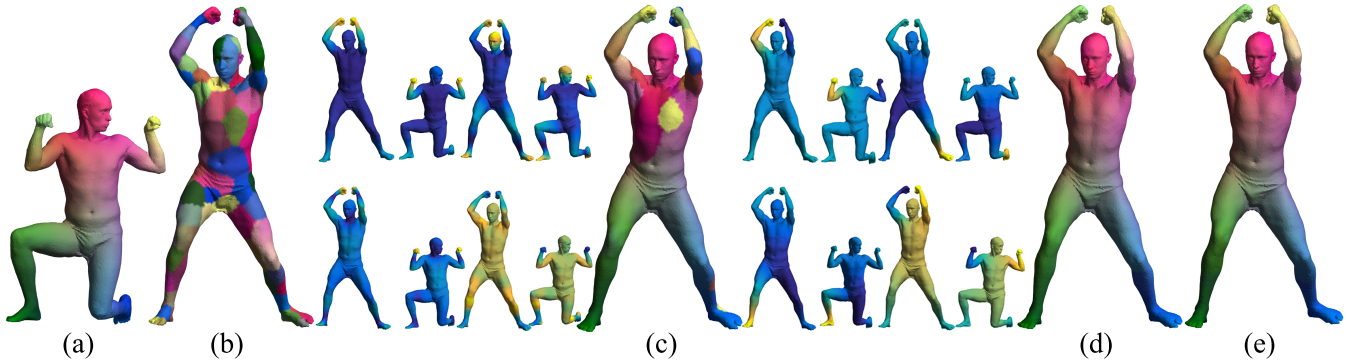

Figure 1: *Initializing high-dimensional ICP for non-rigid registration of two human surfaces using different methods: (a) input shape; (b) random initialization; (c) initialization using Wave Kernel Signatures (four examples are shown to its left); (d) initialization with segment correspondence (four examples are shown to its left). The initialization is crucial for good matching; in (e) we show the result of initialization using PM-SDP which provides comparable result to (d).*

Abstract

Point cloud registration is a fundamental task in computer graphics, and more specifically, in rigid and non-rigid shape matching. The rigid shape matching problem can be formulated as the problem of simultaneously aligning and labelling two point clouds in 3D so that they are as similar as possible. We name this problem the Procrustes matching (PM) problem. The non-rigid shape matching problem can be formulated as a higher dimensional PM problem using the functional maps method. High dimensional PM problems are difficult non-convex problems which currently can only be solved locally using iterative closest point (ICP) algorithms or similar methods. Good initialization is crucial for obtaining a good solution.

We introduce a novel and efficient convex SDP (semidefinite programming) relaxation for the PM problem. The algorithm is guaranteed to return a correct global solution of the problem when matching two isometric shapes which are either asymmetric or bilaterally symmetric.

We show our algorithm gives state of the art results on popular shape matching datasets. We also show that our algorithm gives state of the art results for anatomical classification of shapes. Finally we demonstrate the power of our method in aligning shape collections.

Keywords: Point registration, Shape matching, Convex relaxation

Concepts: •**Computing methodologies** → **Shape analysis;**

SIGGRAPH '16 Technical Paper,, July 24-28, 2016, Anaheim, CA,
ISBN: 978-1-4503-4279-7/16/07

1 Introduction

Registration of point sets is a central problem in computer graphics with many applications including shape analysis, shape retrieval, statistical shape inference, and shape reconstruction.

Among the different formulations of the point set registration problem, the Procrustes matching (PM) formulation is very common: Given two d-dimensional point sets of n points each, $P, Q \in \mathbb{R}^{d\times n}$, which are neither aligned nor consistently labeled, the task is to find a linear isometry (*i.e.*, an orthogonal transformation) $R \in \mathcal{O}(d)$ and a permutation $X \in \Pi_n$ minimizing the distance between the point sets:

$$d(P,Q) = \min_{X,R} \|RP - QX\|_F^2 \tag{1a}$$

$$\text{s.t.} \quad X \in \Pi_n \tag{1b}$$

$$R \in \mathcal{O}(d) \tag{1c}$$

Procrustes matching arises naturally in two and three dimensions ($d = 2, 3$) for rigid matching problems. Non-rigid matching problems are also often formulated in this way, wherein linear isometries in higher dimension ($d \gg 3$) approximate non-rigid isometries of the shapes; this idea is advocated in Functional Maps [Ovsjanikov et al., 2012] where the Laplace-Beltrami eigenfunctions are used for the high-dimensional embedding.

The optimization problem (1) is non-convex and globally optimizing it is difficult. In fact, even the subproblem of finding an exact solution for PM when such a solution exists (*i.e.*, when $d(P,Q) = 0$) is difficult. It can be shown that this subproblem can be solved in polynomial time iff there is a polynomial time algorithm for the exact graph matching problem. The latter is a well researched problem for which no polynomial time algorithm is known.

DOI: http://dx.doi.org/10.1145/2897824.2925913

The iterative closest point (ICP) [Besl and McKay, 1992] algorithm is a popular algorithm for locally minimizing (1), based on the fact that when either R or X are held constant, (1) can be solved globally. As we will demonstrate, (1) has a vast amount of local minima, so that the success of ICP depends heavily on a good initialization.

Previous methods relied on shape features/signatures and/or prior knowledge to initialize ICP. For example, Figure 1 depicts different initializations for solving (1) in the context of Functional Maps for surface matching; a source model (a) is matched to a target model using different initialization for R: (b) shows results achieved from random initialization; (c) demonstrates an initialization of R using Wave Kernel Signatures [Aubry et al., 2011]; and (d) demonstrates initialization using matched segments. Both (c) and (d) are common initializations used in Functional Maps papers [Ovsjanikov et al., 2012; Pokrass et al., 2013]. In this case, all initializations aside from the segmentation correspondence resulted in suboptimal matching (in (c), for instance, the left hand of the model in (a) is incorrectly matched to the chest).

The goal of this work is to approximate the global minimum of (1). To accomplish this, we present a novel convex relaxation of PM using semidefinite programming (SDP), which we name PM-SDP. Standard SDP relaxations are known to give very accurate approximations, at the price of high time complexity. For example, [Kezurer et al., 2015] give an extremely accurate relaxation for the quadratic matching problem, but their algorithm can only run on a handful of points. We formulate a similar SDP relaxation for PM, and use results on semidefinite completion problems to significantly reduce the size of the semidefinite constraints while remaining equivalent to the original relaxation. As a result, our relaxation has significantly improved time complexity.

Our relaxation is applicable to point sets consisting of around 100 points of a reasonably high dimension (15-20). Applying our algorithm to initialize the Functional Maps ICP as shown in Figure 1(e) provides results which compare favorably with matchings achieved with random and WKS initializations, and are comparable to results achieved with matched segmentation initialization.

We demonstrate the accuracy of the suggested relaxation both theoretically and experimentally. We prove that for problems without noise, the relaxation returns a correct global minimum of PM. Up to technical details, this analysis is valid for asymmetric shapes, as well as shapes with bilateral symmetries. The latter include many important instances of the shape matching problem, such as matching human bodies. We also show our algorithm achieves the global optimum for perturbed asymmetric shapes. Experimentally we show that at low noise levels our algorithm still returns the global minimum, and at high noise levels it returns a close-to-optimal local minimum.

We demonstrate the strength and applicability of the PM-SDP relaxation by achieving state of the art results on standard non-rigid shape matching benchmarks such as SCAPE [Anguelov et al., 2005] and the more recent FAUST [Bogo et al., 2014]. We also demonstrate applications to collective matching and biological shape classification.

2 Previous work

Point cloud registration is a basic building block in computer graphics and geometry processing. We mention works that are most relevant to this paper. For detailed survey see [Tam et al., 2013].

For input shapes or point clouds which are close to being aligned, the ICP algorithm and its many variants [Rusinkiewicz and Levoy, 2001] are a popular choice. For shapes that are not roughly aligned the ICP algorithm can be easily stuck in local minimum. A widely used method for matching shapes in a general position is RANSAC [Fischler and Bolles, 1981], which exhaustively samples the space of possible transformations. Being essentially a brute-force algorithm, its drawback is the high complexity which results in poor scalability to higher dimensions.

Another way to deal with shapes that are not roughly aligned is to try to find a good initialization for ICP. Representative works in this direction are the works of [Gelfand et al., 2005; Yang et al., 2013] which use combinatorial optimization techniques in order to find the optimal solution in 3D.

Non-rigid shape matching is a central task in geometry processing. We discuss relevant work and refer to relevant surveys for more details [Wand et al., 2011; Van Kaick et al., 2011; Tam et al., 2013]. One approach is to formulate the non-rigid problem as a version of the quadratic assignment problem called quadratic matching. [Mémoli and Sapiro, 2004; Bronstein et al., 2006] suggest to minimize the Gromov-Hausdorff distance where [Bronstein et al., 2006] also introduce the generalized multidimensional scaling variant on point metric spaces; [Leordeanu and Hebert, 2005; Berg et al., 2005] relax the quadratic matching problem using linear programming and spectral techniques; [Kezurer et al., 2015] suggest a convex SDP relaxation to the quadratic matching problem; and [Chen and Koltun, 2015] suggest a linear programming relaxation solved using MRF (Markov Random Fields) techniques.

Another approach is to restrict the mapping search space to a smaller, more tractable space such as: low dimensional deformation space [Brown and Rusinkiewicz, 2007], conformal mappings [Lipman and Funkhouser, 2009; Zeng et al., 2010; Kim et al., 2011]; or isometries [Ovsjanikov et al., 2010; Tevs et al., 2009]. Some works try to adapt the 3D ICP algorithm to the non-rigid case [Li et al., 2008], whereby typically a deformation model is chosen (*e.g.*, piecewise affine) and the algorithm iterates between finding correspondences and solving for the optimal deformation.

Non-rigid shape matching can also be tackled using supervised machine learning techniques, where typically the algorithm is focused on a specific type of data. For instance, [Zuffi and Black, 2015; Wei et al., 2015] train a model that matches human bodies.

The most relevant methods to this paper pose the non-rigid shape matching problem as a high dimensional rigid shape matching problem [Jain et al., 2007; Ovsjanikov et al., 2008; Ovsjanikov et al., 2012; Pokrass et al., 2013]. These methods map each point of the input point cloud to a vector containing the values of the eigenfunctions of the Laplace-Beltrami operator [Rustamov, 2007]. The central observation is that if the input shapes are isometric, the intrinsic isometry between them becomes an extrinsic isometry in the high dimensional space [Ovsjanikov et al., 2008]. This opens the door for using rigid matching algorithms such as ICP in the context of the more complicated non-rigid matching problem. The problem is that finding a good initialization is difficult for high dimensional problems. Current methods partly employ prior knowledge, such as correspondence between pre-computed segments, to initialize the ICP algorithm [Ovsjanikov et al., 2012].

Procrustes analysis is a tool used to perform statistical study of shapes by canceling transformations that are not shape-altering [Kendall, 1984]. It has many applications in various fields of science [Gower and Dijksterhuis, 2004] and in particular anatomical shape analysis and morphometrics [Mitteroecker and Gunz, 2009; Boyer et al., 2011]. There are many variants to the Procrustes problem, maybe the most general is the one addressed in this paper - PM, for which no closed-form solution is known [Gower and Dijksterhuis, 2004]. To our knowledge, we provide the first closed-form convex formulation guaranteeing a globally optimal solution for the exact and near-exact cases (under mild assumptions).

Semidefinite relaxation and polynomial optimization. Convex relaxations of quadratic optimization problems are often (*e.g.*, [Poljak et al., 1995; Luo et al., 2010]) preformed by replacing quadratic terms with new variables. All quadratic terms then become linear, and a convex semidefinite constraint is added to couple between the original variables and the new variables. A significant drawback of these relaxations is that they are computationally tractable only for very small polynomial problems. [Fukuda et al., 2000; Waki et al., 2006] show that for problems with certain structure, a positive semidefinite constraint on a large matrix can be replaced by constraining certain principal submatrices to be positive semidefinite, resulting in an equivalent problem with significantly improved time complexity. In this paper we devise a quadratic formulation of PM which has this structure, and as a result obtain a relaxation which is tractable for medium sized problems.

3 Approach and Formulation

We present our convex relaxation of PM. We first discuss the general quadratic optimization problem and present a novel strategy for replacing its "standard" time consuming SDP relaxations with an equivalent, but significantly more efficient SDP relaxation. In this context, we then instantiate our relaxation for the PM problem.

Full SDP relaxation of quadratic problems. Quadratic optimization problems are problems of the form

$$\min_{x\in\mathbb{R}^N} \quad f_0(x) \tag{2a}$$

$$\text{s.t.} \quad f_s(x) = 0, \quad s = 1,\dots,S \tag{2b}$$

$$f_t(x) \geq 0, \quad t = S+1\dots T \tag{2c}$$

where f_i are quadratic multivariate polynomials.

A typical relaxation procedure includes two steps (see [Luo et al., 2010] for a survey on SDP quadratic relaxations): First, the quadratic polynomials f_i are linearized by introducing new variables Y_{ij}, $1 \leq i,j \leq N$, which replace quadratic monomials $x_i x_j$, so that $f_j(x)$ becomes a linear polynomial in the variables x, Y denoted $\mathcal{L}[f_j](x,Y)$. This gives an equivalent formulation of (2):

$$\min_{x,Y} \quad \mathcal{L}\left[f_0\right](x,Y) \tag{3a}$$

$$\text{s.t.} \quad \mathcal{L}\left[f_s\right](x,Y) = 0, \quad s = 1\dots S \tag{3b}$$

$$\mathcal{L}\left[f_t\right](x,Y) \geq 0, \quad t = S+1\dots T \tag{3c}$$

$$Y = xx^T \tag{3d}$$

With the exception of constraint (3d), Problem (3) is a convex problem (in fact a linear program). Therefore, the second step in this relaxation procedure is replacing (3d) with a convex constraint. The convex hull of the set defined by (3d) is the set defined by the convex constraint $Y \succeq xx^T$, which is equivalent to the semidefinite constraint

$$\begin{bmatrix} 1 & x^T \\ x & Y \end{bmatrix} \succeq 0. \tag{4}$$

A natural relaxation of (3) is therefore given by replacing (3d) with (4). [Zhao et al., 1998] and more recently [Kezurer et al., 2015] used this approach to relax the quadratic matching and quadratic assignment problems. The obtained relaxation is significantly more accurate than prevalent relaxations for quadratic matching, but its scalability is poor; in fact, it cannot handle more than a handful of points to be matched, completely hindering some applications.

Efficient SDP relaxation. The key to a useful and efficient relaxation of (2) and consequently our problem (1) is to reduce the dimension of the semi-definite constraint (4) which is the main factor determining time efficiency of the semidefinite program.

To obtain a more efficient SDP relaxation, we make the observation that for some problems not all terms in the matrix xx^T appear in the polynomials f_j. This is, for example, the case in the PM problem, as will soon be shown. In such cases, we can find a collection $\mathcal{J}$ of subsets of $\{1,\dots,N\}$ so that all polynomials f_j include only expressions from $x_J x_J^T,\ J \in \mathcal{J}$. An equivalent formulation for (3) can therefore be obtained by replacing (3d) with $Y_J = x_J x_J^T$, for all $J \in \mathcal{J}$. In turn, replacing these with the convex constraints

$$\begin{bmatrix} 1 & x_J^T \\ x_J & Y_J \end{bmatrix} \succeq 0, \qquad J \in \mathcal{J} \tag{5}$$

we obtain a convex relaxation for (2). If all subsets $J \in \mathcal{J}$ satisfy $|J| \ll N$, the obtained relaxation is considerably more efficient than the original (full) relaxation.

There is no unique way to apply this more efficient relaxation; a given instance of a quadratic optimization problem may have several different possible decompositions $\mathcal{J}$. Those that use small semidefinite constraints will be more efficient, but not necessarily as accurate as those using larger semidefinite constraints; the latter, however, can quickly become intractable for certain problems. Nevertheless, if $\mathcal{J}$ is chosen so that it satisfies the chordality condition we will soon describe, the obtained relaxation is in fact equivalent to the full relaxation.

In general, any solution for the full relaxation also satisfies (5). For equivalency, we need to ensure that a solution for the efficient relaxation (5) can always be completed to a solution of the full relaxation. For that end we need to show there is a solution for the following matrix completion problem: We are given entries of x_J, Y_J satisfying (5), and we are searching for a completion of Y that satisfies (4). Since the objective and linear constraints depend only on the coordinates which were determined before the completion, the full solution will also fulfill the linear constraints, and the objective will not be affected by the completion.

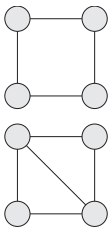

The condition that allows solving the completion problem is related to the structure of the known coordinates of the matrix. The collection $\mathcal{J}$ defines an undirected graph $G = (V,E)$ whose vertices are $V = \{1, x_1, \dots, x_N\}$. Two distinct vertices are connected by an edge iff they both appear in one of the matrices (5) defined by some $J \in \mathcal{J}$. A graph G is *chordal* if every (simple) cycle with more than three vertices contains a chord, *i.e.*, an edge between two non-adjacent members of the cycle. For example the graph in the top of the inset has a cycle which does not contain a chord and thus is not a chordal graph. The bottom graph is chordal.

If G is chordal, the following theorem from [Grone et al., 1984] guarantees that the matrix completion problem has a solution, and therefore that the two relaxations are equivalent:

Theorem 1. *If G is chordal, and $(x_J, Y_J)_{J\in\mathcal{J}}$ satisfy* (5)*, then the missing coordinates of Y can be chosen so that the full semidefinite constraint* (4) *holds.*

PM-SDP Formulation. We now return to the PM problem and instantiate the strategy presented above. First, we note that PM can be formulated as the following quadratic problem:

$$\min_{X,R} \|RP - QX\|_F^2 \tag{6a}$$

$$X\mathbf{1} = \mathbf{1}, \quad \mathbf{1}^T X = \mathbf{1}^T \tag{6b}$$

$$X_j X_j^T = \operatorname{diag}(X_j), \quad j = 1\dots n \tag{6c}$$

$$RR^T = R^T R = I \tag{6d}$$

where we denote by $\mathbf{1} \in \mathbb{R}^{n\times 1}$ the all-ones vector, by X_j the j-th column of X, and by $\mathrm{diag}(X_j)$ the diagonal matrix whose diagonal entries are X_j.

To see this is indeed an equivalent formulation of PM, note that if (R, X) is a feasible solution of (6), then R is orthogonal by definition. The constraint (6c) implies that $X_{ij}^2 = X_{ij}$ for all i, j, so that all elements of X are in $\{0, 1\}$. By (6b) the rows and columns of X sum to one, which implies that X is a permutation matrix.

In the other direction, note that if $(R, X) \in \mathcal{O}(d) \times \Pi_n$, then since each column of X has only one non-zero element, $X_j X_j^T$ is diagonal. Since all elements of X are in $\{0, 1\}$, $X_{ij}^2 = X_{ij}$ for all i, j so that (6c) holds. It is straightforward to check that (6b),(6d) hold as well.

All polynomials in (6) are quadratic in the entries of R and X. The full SDP relaxation for quadratic problems described above can then be applied; this will result in a vector x, consisting of the elements of R and X, of dimension $d^2 + n^2$; the SDP constraint will be of size $(d^2 + n^2 + 1) \times (d^2 + n^2 + 1)$.

We obtain an equivalent, more efficient, relaxation by utilizing the efficient SDP relaxation approach; the key observation here is that all the quadratic polynomials participating in the formulation (6) of the PM problem can be expressed using linear polynomials in the entries of $X_j X_j^T$, $X_j [R]^T$, and $[R] [R]^T$, where $[R] \in \mathbb{R}^{d^2\times 1}$ is the column stack of the matrix R and $j = 1 \dots n$. We therefore introduce new matrix variables Z_j, constrained to satisfy

$$Z_j = \begin{bmatrix} X_j \\ [R] \end{bmatrix} \begin{bmatrix} X_j \\ [R] \end{bmatrix}^T, \quad j = 1 \dots n \tag{7}$$

Let us next see how the objective and constraints of Problem (6) are linear in the variables X, R, Z_j. First, the objective of (6) can be rewritten as

$$\|RP - QX\|_F^2 = \sum_j \|RP_j - QX_j\|_2^2 = \sum_j \mathrm{tr}\,(W_j Z_j) + const$$

for some constant matrices W_j since $\|RP_j - QX_j\|_2^2$ is linear in the entries of Z_j. Denoting

$$Z_j = \begin{bmatrix} A_j & B_j^T \\ B_j & C \end{bmatrix},\ A_j \in \mathbb{R}^{n\times n}, C \in \mathbb{R}^{d^2\times d^2}, B_j \in \mathbb{R}^{d^2\times n} \tag{8}$$

the constraint (6c) can be rewritten as $A_j = \mathrm{diag}(X_j)$. Finally, the constraints (6d) are affine functions of C and can be rewritten as

$$\mathrm{tr}(H_\ell C) + b_\ell = 0, \quad \ell = 1 \dots 2d^2$$

for some constant matrices H_ℓ. Replacing the non-convex equality constraint of (7) with convex semidefinite constraints of type (4), we obtain our relaxation for the PM problem, PM-SDP:

$$\min_{Z_j, X, R} \quad \sum_j \mathrm{tr}\,(W_j Z_j) \tag{9a}$$

$$X\mathbf{1} = \mathbf{1}, \quad \mathbf{1}^T X = \mathbf{1}^T \tag{9b}$$

$$A_j = \mathrm{diag}(X_j), \qquad j = 1 \dots n \tag{9c}$$

$$\mathrm{tr}(H_\ell C) + b_\ell = 0, \qquad \ell = 1 \dots 2d^2 \tag{9d}$$

$$Z_j \succeq \begin{bmatrix} X_j \\ [R] \end{bmatrix} \begin{bmatrix} X_j \\ [R] \end{bmatrix}^T, \qquad j = 1 \dots n \tag{9e}$$

where A_j and C are defined as in (8).

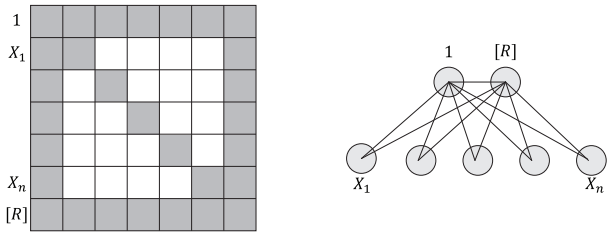

Figure 2: *The graph corresponding to the Procrustes Problem (right, each disk represents a clique) is chordal, that is, has no minimal cycles of length at-least four. The adjacency matrix is shown on the left.*

Chordality of the relaxation. To show that (9) is equivalent to the full relaxation of the PM problem, we need to show that the graph G induced by $\mathcal{J} = \{J_j\}$ is chordal, where J_j is the set containing the variables $[R]$ and X_j. The adjacency matrix of G for this case is illustrated in Figure 2 (left) where each gray square represents a full block of ones; on the right, we illustrate the corresponding graph G where each disk represents a clique which corresponds to a diagonal block in the adjacency matrix. To show G is chordal it is enough to show every cycle of length at-least 4 has a chord: Indeed, if a cycle is completely contained in one of the sets J_j (represented as triangles in Figure 2, right) then any two vertices are connected by an edge; otherwise there are two non consecutive visits to vertices in the $1, [R]$ cliques (top disks), and these vertices are connected by an edge.

Dimension and complexity. We note that (9e) includes n semidefinite constraints involving $N \times N$ matrices, where $N = d^2 + n + 1$, as opposed to the full relaxation which in this case would involve a square matrix with $N = d^2 + n^2 + 1$. When the number of points in the sets n is significantly larger than the dimension d of the space the point reside in (*i.e.*, $n \gg d$), which is often the case in point registration problems, PM-SDP will be significantly more efficient than the full SDP relaxation.

Relaxation properties. The PM-SDP relaxation has the following natural theoretical properties:

1. *Rotation and relabeling invariance*: Rotating or relabeling the input shapes will not affect the solution provided by the relaxation. More precisely, if P is replaced with a point cloud $\tilde{P}$ obtained from P by relabeling and applying an orthogonal transformation, then the objective of PM-SDP remains unchanged, and X, R transform accordingly.
2. *Lower bound*: Since PM-SDP is a relaxation its optimal objective is less or equal to the PM optimal objective. We denote the objective of PM-SDP by $\underline{d}$.
3. *Positivity*: The objective of PM-SDP is always non-negative. This is natural since this is the case for the objective of PM .
4. *Convex-hull of R, X*: The R, X coordinates of a feasible solution for PM-SDP are in the convex hull of $\mathcal{O}(d) \times \Pi_n$.

The proof of the these properties as well as the theorems presented below are given in [Dym and Lipman, 2016].

Exact recovery refers to the problem of finding $R, X \in \mathcal{O}(d) \times \Pi_n$ which solve the equation $RP = QX$, when such solutions exists, *i.e.*, when $d(P, Q) = 0$. We call such solutions *exact solutions*. From the computational perspective, exact recovery for PM is

equivalent to exact graph matching. The latter is a well-researched problem, not known to be polynomial. Accordingly, proving exact recovery in full generality is not likely. However, under the assumption that the covariance $d \times d$ matrix PP^T of the point cloud P has a simple spectrum, and an additional weak assumption, we are able to prove exact recovery. This too is analogous to the graph matching problem where finding exact solutions for graphs with simple spectrum affinity matrices is solvable in polynomial time [Babai et al., 1982].

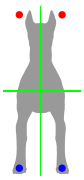

The assumption that PP^T has a simple spectrum implies that the symmetries of P are all reflections along the principle axes of the point set P. In particular, it implies that all symmetries of P are bilateral. The class of bilateral symmetric shapes includes many important instances of the shape matching problem. Note however, that not all reflections along principle axes are necessarily symmetries of P. The additional weak assumption required for our exact recovery with symmetries result formulated below is that there exists a point P_j in P such that its reflections along principle axes belongs to P only for symmetries of P. The inset figure demonstrates a shape with point having this property (blue); Applying a horizontal flip, which is a symmetry, maps the point to another point on the shape, while applying a (non-symmetry) vertical flip maps it outside of the shape (red). We are not aware of bilateral symmetric shapes of practical interest that do not satisfy this condition.

The exactness argument starts with assuming we are given P, Q with $d(P,Q) = 0$ and showing that when $d(P,Q) = 0$ also $\underline{d}(P,Q) = 0$. This follows from relaxation properties 2 and 3 described above: from the lower bound property we know that the objective $\underline{d}$ of PM-SDP is a lower bound of the objective d of PM. From the positivity property, $\underline{d}$ is always non-negative. It follows that when $d(P,Q) = 0$ also $\underline{d}(P,Q) = 0$, and the set of feasible solutions of PM-SDP with zero objective, which we call the *exact convex solution set*, is a superset of the set of exact solutions. When the shapes are asymmetric, the exact convex solution set consists of only one point - the exact solution,

Theorem 2. *Let P, Q be asymmetric shapes with $d(P,Q) = 0$ satisfying the simple spectrum and weak conditions. Then PM-SDP has a unique exact convex solution, which is also the unique exact solution of PM.*

When P, Q are bilateral symmetric, there are several exact solutions. All convex combinations of these solutions will be in the exact convex solution set, so that generally exact convex solutions will not be exact solutions. However, by restricting ourselves to the R coordinate of both the exact solutions and the convex exact solutions, which we refer to as *exact orthogonal solutions* and *exact convex orthogonal solutions*, we are able to show:

Theorem (Full version of theorem 2)**.** *Let P, Q be shapes with $d(P,Q) = 0$ satisfying the simple spectrum and weak conditions. Then the exact orthogonal solutions of PM are the extreme points of the set of exact convex orthogonal solutions.*

The set of exact convex orthogonal solutions is a convex set and therefore its extreme points can be found by simply optimizing linear energies over this set (a convex problem again). One simple algorithm for obtaining *all* exact solutions is repeatedly solving a variation of our convex relaxation: minimize a random linear energy $\mathrm{tr}(WR)$, where $W \in \mathbb{R}^{d \times d}$ is a random matrix drawn from the uniform measure on the unit sphere, under the set of constraints of (9), and adding the linear constraint that the objective (9a) is zero. We prove,

Theorem 3. *The random algorithm returns an extreme point of the set of exact convex orthogonal solutions (i.e., an exact orthogonal solution) with probability one. Moreover, all extreme points are found with the same probability.*

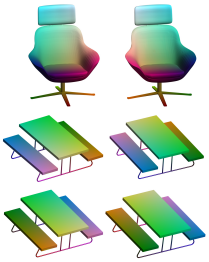

Once an exact orthogonal solution is found it can be shown that the X coordinate of the solution is guaranteed to be a permutation. The inset demonstrates the output of the probabilistic algorithm described in theorem 3 on two point sets with perfect bilateral symmetries sampled from a model of a chair and a picnic table. The random algorithm retrieves the two bilateral symmetries of the chair, and the four bilateral symmetries of the picnic table.

4 Implementation details

We discuss implementation details of the PM-SDP algorithm.

Injective matching. We consider a slight variation of PM where we allow the point cloud P to have fewer points than Q, and search for the correspondence between P and a subset of the points of Q. This formulation is useful to account for the inherent noise caused from samplings of different shapes. Furthermore, in case the shapes we wish to compare are not isometric, certain points on one shape might not have a good match on the second one.

We denote the number of points of P by $k \leq n$. The formulation of PM in (1) remains unchanged, except now $X \in \mathbb{R}^{n \times k}$ is constrained to be a matrix with entries in $\{0, 1\}$ such that all columns of X have exactly one non-zero entry. The only necessary modification for PM-SDP is that the constraint $X\mathbf{1} = \mathbf{1}$ in (9) should be replaced with the constraint $X\mathbf{1} \leq \mathbf{1}$.

Utilizing priors for computational efficiency. The PM-SDP framework allows incorporating priors to improve computational complexity by further reducing the size of the SDP constraint. This is done by noting that if we know or assume the points P_j and Q_i should not correspond then setting $X_{ij} = 0$ reduces the size of the SDP constraint by one. Indeed, returning to (7) we see that when $X_{ij} = 0$ we may also assume that the i-th row and column of Z_j are also zero. Similarly, if $R_{st} = 0$ we can eliminate a row and column of Z_j.

Let us demonstrate where this can be used for near-isometric matching. We rule out unlikely correspondences using the average geodesic distance (AGD) descriptor [Kim et al., 2011]. Corresponding points on nearly isometric shapes should have roughly the same AGD value since this descriptor is isometry-invariant. For close to isometric shapes, we can safely rule out correspondences between points whose AGD value is significantly different. If the possibility of Q_i corresponding to P_j, is ruled out, we set $X_{ij} = 0$.

For isometric shapes, the Laplace-Beltrami operator of both shapes has the same eigenvalues, and the linear isometry R takes eigenfunctions of the first shape with eigenvalue λ to eigenfunctions of the second shape with the same eigenvalue λ. In the most common simple spectrum case, this means that R is a diagonal matrix, if the maximal eigenvalue multiplicity of the shapes is two then R is tridiagonal, etc. For near-isometric shapes, we make the assumption that R is m-diagonal, that is, has m non-zero diagonals symmetrically around the main diagonal.

In practice, we use the AGD descriptor to rule-out 50-70% of the correspondences, and constrain R to be 1,3, or 5-diagonal. Fortunately, the effect of incorporating these priors on the quality of the relaxation is negligible: Figure 3 demonstrates that using PM-SDP to match shapes from SCAPE dataset (using the protocol that will be described in Section 6) produces essentially equivalent results with and without using the AGD descriptor to rule out unlikely matches; the former, however, has the advantage of significantly improved computational efficiency.

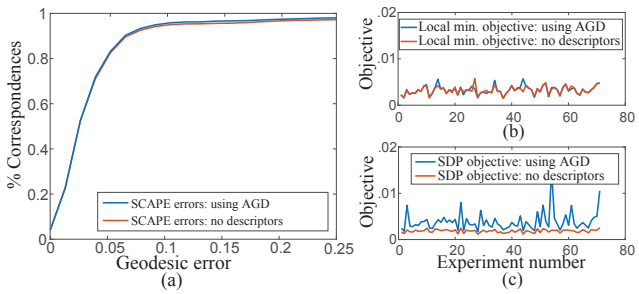

Figure 3: *Ruling out matches with the AGD descriptor has a negligible effect on the quality of the relaxation: (a) depicts the results of PM-SDP on SCAPE dataset [Anguelov et al., 2005] with and without AGD to rule out unlikely matches; (b) the objective after local minimization; and (c) objective value achieved by PM-SDP. The PM-SDP objective is lower for the unpruned version while the rest of the results are equivalent for both versions.*

Local minimization. Since the feasible set of PM-SDP is larger than the feasible set of PM, the solution of PM-SDP in general may not contain orthogonal and permutation matrices. We therefore project the solution onto the feasible set of PM. We do this by locally minimizing PM using the output of PM-SDP to initialize the algorithm. The local minimization is done using an ICP-like algorithm which interleaves between minimizing over one of the matrices R, X while holding the other constant: fixing R results in a linear program, while for a fixed X there exists a closed-form solution [Gower and Dijksterhuis, 2004]. In Figure 4 we illustrate the doubly-stochastic matrix X as outputted from the PM-SDP relaxation and the permutation achieved after the projection. As shown, the PM-SDP output is already very similar to the projection result demonstrating the tightness of the PM-SDP relaxation. More details are in Appendix A.

The local minimization following the PM-SDP relaxation allows generalizing Theorem 2 to the inexact case:

Corollary 1. *Let P, Q be point clouds satisfying the conditions of Theorem 2, and let P^ϵ, Q^ϵ be sufficiently small perturbations of P, Q. Then PM-SDP followed by the local minimization returns the unique (global) solution of PM for P^ϵ, Q^ϵ.*

5 Evaluation

We test the tightness of the PM-SDP relaxation by comparing it to the ground truth obtained from an exhaustive brute-force sampling algorithm. The latter is only tractable for low dimensional d, and we choose $d = 3$: The exhaustive algorithm densely samples $\sim 10k$ points from a uniform distribution over $\mathcal{O}(3)$ and uses each sample R_j as an initialization for the local minimization algorithm described above.

In Figure 5 we compare the histograms of optimal values achieved by the exhaustive sampling algorithm (in red) to the energy achieved by PM-SDP (in blue). The data for this experiment was created by randomizing $Q \in \mathbb{R}^{3\times 50}$ according to a uniform distribution on $[0, 1]$, and setting $P = R^T QX + \varepsilon$, with $X \in \Pi_{50}$, $R \in \mathcal{O}(3)$ and noise $\varepsilon \sim N_{d\times n}(0, \sigma^2)$. (a-d) show the results of a few typical runs with increasing amount of noise $\sigma = 0, 0.05, 0.1, 0.2$. We note that the number of local (sub-optimal) minima for the exhaustive sampling is surprisingly high; for example, for noise level $\sigma = 0.1$ we found more than 1000 local energy minima. Additionally, the experiment in (a) verifies our theoretical exactness result as can be seen by the fact that the blue point achieves the left most value of the red histogram. When the noise level is low to medium ($\sigma = 0.05, 0.1$) the PM-SDP relaxation usually produces optimal result, see (b-c). When noise level is high ($\sigma = 0.2$ in (d)) the relaxation does not provide an optimal solution but nevertheless produces a close to optimal result.

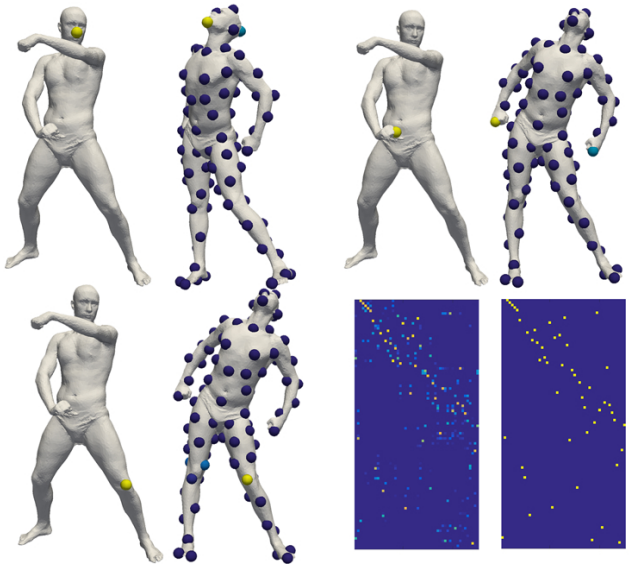

Figure 4: *Visualization of the doubly-stochastic map X as generated by the PM-SDP relaxation when comparing two SCAPE models; each pair of surfaces depicts a column of X by coloring the point set Q according to the corresponding value in X; the X matrix before and after projection on the permutations is shown at the bottom-right.*

A quantitative evaluation of the optimality of PM-SDP is given in Figure 5, (e-f). We ran 80 random experiments for $d = 3$ and $d = 5$ with noise level $\sigma = 0.1$ and measured the optimal objective value achieved by PM-SDP in comparison to the global minimum and median value of the objective values found by the exhaustive algorithm. For visualization, we subtracted the value of the optimal value from all of the results. PM-SDP (black line) usually returns the optimal value (green) and always returns a better result than the median objective value (blue) of the exhaustive algorithm.

6 Applications

6.1 Functional maps

The main application of our algorithm is non-rigid shape matching of pairs of surfaces $\mathcal{P}, \mathcal{Q}$. Following [Ovsjanikov et al., 2012] we pose this problem as a high dimensional PM problem, replacing non-rigid isometries with linear isometries (orthogonal transformations) in higher dimensional space. More specifically, we sample k points on the first shape and n points on the second shape uniformly using farthest point sampling [Eldar et al., 1997] initialized with extrema of average geodesic distance (AGD) [Kim et al., 2011], and embed P, Q in $\mathbb{R}^d$. The embedding is done by first computing the first d eigenfunctions of the cot-weight Laplace-Beltrami (LB) operator [Pinkall and Polthier, 1993] on each of the surfaces, $\{\Phi_i^{\mathcal{P}}\}$, $\{\Phi_i^{\mathcal{Q}}\}$, $i = 1 \dots d$ and then assigning the d coordinates $(\Phi_1^{\mathcal{P}}(p), .., \Phi_d^{\mathcal{P}}(p))$ to each point $p \in P$, and similarly to every point $q \in Q$.

Current approaches using this formulation, solve the resulting high dimensional PM problem using an ICP-type iterative algorithm; as this problem is shown to have a vast number of local minima even for $d = 3$ (see Figure 5 (a-d)), initialization is crucial.

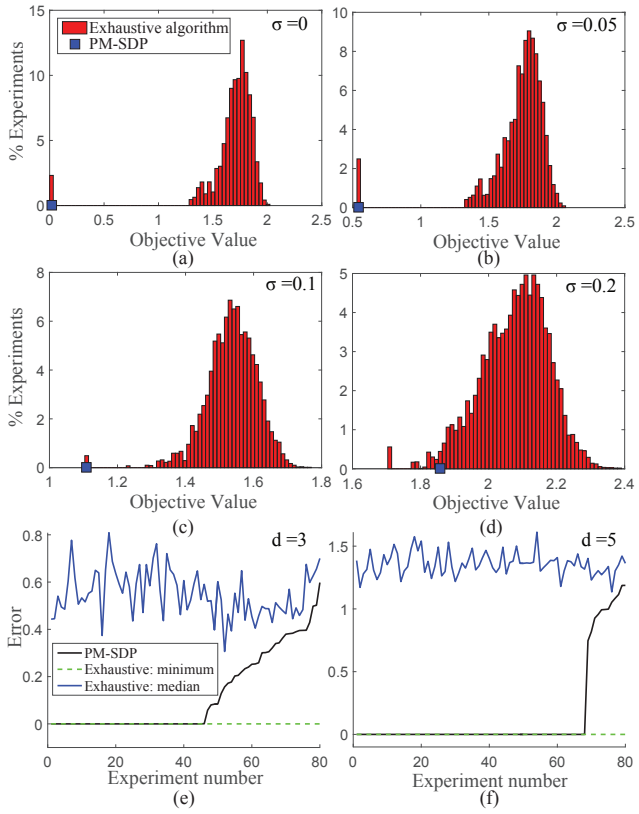

Figure 5: *PM-SDP tightness evaluation: (a-d) show the histograms of the objective values achieved by the exhaustive sampling algorithm compared to the optimal PM-SDP objective on a few typical runs. When the noise level is low to medium, our algorithm usually finds the global minimum. On higher noise levels it returns an objective value close to optimal. (e-f) show illustrations of the deviation of the optimal PM-SDP objective value from the global optimum and the median value computed by the exhaustive algorithm when $\sigma = 0.1, d = 3$ and $d = 5$.*

Using standard shape signatures or features often does not provide a satisfactory initialization (*e.g.*, Figure 1, (b-c)), and previously this ICP procedure was initialized with matched segments for successful results [Ovsjanikov et al., 2012; Pokrass et al., 2013]. In the experiments of this subsection we use PM-SDP followed by local minimization to initialize the ICP of [Ovsjanikov et al., 2012] that uses the entire embedded models. In comparisons to previous works we used code supplied by the authors with their default parameters.

SCAPE dataset. We evaluated the performance of PM-SDP for non-rigid isometric matching using the SCAPE dataset [Anguelov et al., 2005]. We used $n = 100$ sampled points, $d = 17$ eigenfunctions, and injective mapping of $k = 50$ out of $n = 100$. For better efficiency we allowed each point in P to be matched only to the 30% of the points in Q that have the closest AGD, and selected $m = 5$ (5 non-zero diagonals in R). The SDP optimization was performed using Mosek [MOSEK, 2015]. We extended our results to a full correspondence of all the vertices by solving ICP in dimension $d = 30$. The average running time for a pair with these settings is 30-35 minutes on an Intel Xeon E5 CPU. We also tested a faster version by taking $d = 30$ and using diagonal R, that is $m = 1$; this gave only slightly inferior results with running time of 2.5 minutes per pair.

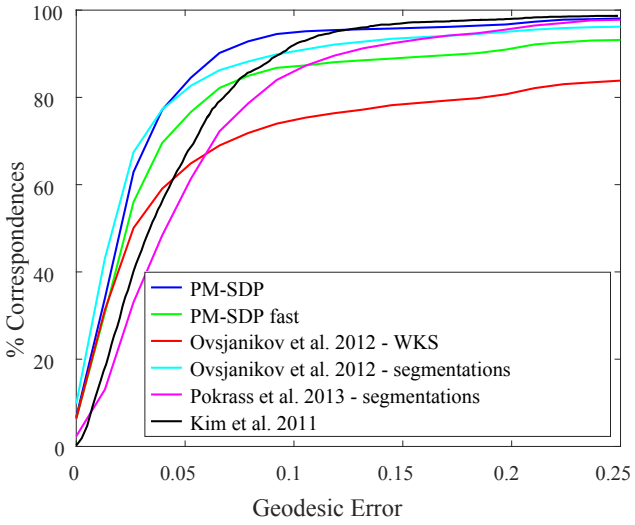

Figure 6: *Results of our algorithm and state of the art algorithms on the SCAPE [Anguelov et al., 2005] non-rigid shape matching benchmark.*

Figure 6 shows a comparison of our algorithm with several state of the art algorithms. The comparison was done according to the protocol of [Kim et al., 2011] accepting symmetries. Our method compares favorably to functional maps (FM) when initialized with matched segments [Ovsjanikov et al., 2012] and improves upon FM with automatic shape signature initialization. We also show in green the results of the faster, less accurate, variant of our algorithm described in the previous paragraph. Figure 7 shows typical results of our algorithm from this experiment.

FAUST dataset. We evaluated the performance of PM-SDP for non-rigid non-isometric matching on the FAUST dataset [Bogo et al., 2014]. We used a similar setup as in the previous experiment, with the following differences: We generated the LB operator directly on the point cloud sampling generated by [Chen and Koltun, 2015] using a similar construction to [Belkin and Niyogi, 2003] (weights based on geodesic distances instead of Euclidean distances). We also used two versions of the PM-SDP: For the first we chose $d = 17$, $m = 5$ and we allowed each point in P to be matched only to the 40% points in Q that have the closest AGD; the running time of this parameters set is about 40 minutes per pair. For the second, faster parameter set, we used $n = 40, k = 30$, an embedding with $d = 10$, $m = 5$, and kept 80% of Q for each point in P; the running time with these parameters is less than 4 minutes per pair.

Figure 8 compares PM-SDP with the recent method of [Chen and Koltun, 2015] which demonstrated superb state of the art results on this dataset. However, they rely on the assumption that the shapes are initially aligned in 3D and indeed use this alignment by adding a regularization term. In order to make a fair comparison we disabled this regularization term. When this term is removed, intrinsic symmetries might be found by the algorithm. In order to account for that we sampled a set of 52 ground truth points in each mesh, and added the symmetric flip to the ground truth map. Aside from that, we followed Chen and Koltun's evaluation protocol (including using their point clouds as stated above). As can be read from the graphs, our algorithm (blue) compares favorably in both the inter and intra class matching scenarios in terms of cumulative error distribution and average error. We also show here a faster version of our algorithm (green), which provides good results in a shorter computation time. Figure 9 shows some typical results of our algorithm for both inter and intra class matching.

Figure 7: *Examples of typical maps obtained with PM-SDP on the SCAPE dataset [Anguelov et al., 2005]. In all pairs: left mesh is colored using a predefined color map; right mesh is colored according to the correspondence. Bottom right: a failure case.*

SCAPE dataset (raw scans). We further tested our algorithm on the SCAPE original raw scans dataset [Anguelov et al., 2005] that contain missing data, holes and noise. We used the same preprocessing method of [Chen and Koltun, 2015] and ran our algorithm with exactly the same parameters as on the FAUST dataset on the 71 pairs as defined in the benchmark of [Kim et al., 2011]. Figure 10 shows the cumulative error graph and a few typical results. We note that also here we ran [Chen and Koltun, 2015] without the extrinsic regularization term (in addition to the reasons stated above, for the SCAPE dataset this prior is inappropriate due to its pose diversity).

SHREC07 dataset. We also ran PM-SDP ($n = 100, k = 40$) on the highly non-isometric SHREC07 dataset [Giorgi et al., 2007]. On this dataset, PM-SDP achieved good results only on some of the classes; Figure 11 demonstrates typical results from these classes: the Ant, Teddy and Glasses.

6.2 Anatomical classification

The Procrustes distance with labeled points (*i.e.*, when X is known) is a well-known measure of shape similarity in fields such as statistical shape analysis [Kendall, 1984; Boyer et al., 2011]. The sampling and labeling of points in a collection of shapes is tedious work that requires the attention of an expert for several months [Boyer et al., 2011]. The possibility of solving Procrustes matching with unlabeled points (*i.e.*, the PM problem in this paper) using PM-SDP makes the task of finding meaningful landmarks unnecessary.

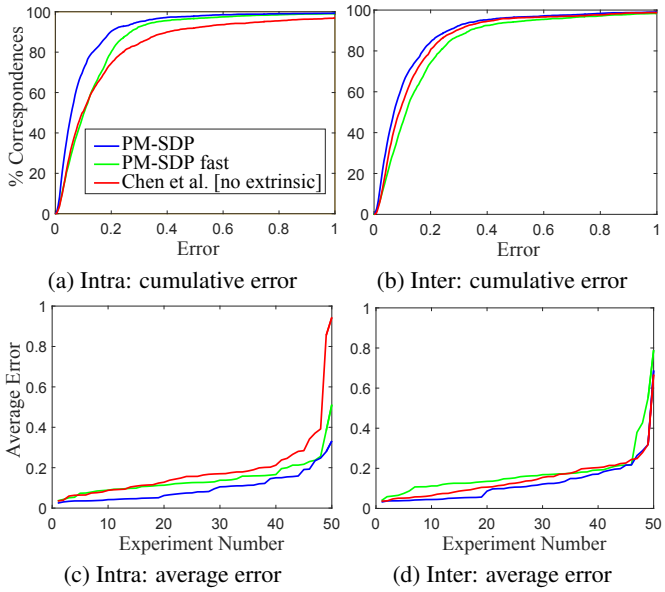

(a) Intra: cumulative error (b) Inter: cumulative error

(c) Intra: average error (d) Inter: average error

Figure 8: *Cumulative and average errors achieved on the FAUST dataset [Bogo et al., 2014] by PM-SDP compared to [Chen and Koltun, 2015] without the global extrinsic regularization term.*

We took three anatomical bone datasets containing 116, 61 and 45 models respectively from [Boyer et al., 2011]. We sampled $n = k = 120$ points of each shape using farthest point sampling, ran PM-SDP and used its output to initialize ICP that matches 400 farthest points on the shapes. This computation takes about 7 minutes for each pair.

We followed the classification protocol suggested in [Boyer et al., 2011] where each shape is classified according to its nearest (in Procrustes distance) neighbor; each shape in the datasets has three biological tags: Genera, Family and Above Family, and we tested classification of all three categories. Table 1 presents classification success rates (what fraction of shapes were correctly classified in each classification test) and shows PM-SDP compares favorably to Boyer's method [Boyer et al., 2011], and is remarkably comparable to the results achieved using human expert labeled landmarks. Figure 12 shows a few examples of maps that were found by PM-SDP.

6.3 Shape collection alignment

We demonstrate another application of PM-SDP to consistent alignment of shapes. The task we would like to solve here is the following: given a set of semantically similar shapes - apply an orthogonal transformation per shape so that the shapes are aligned. We solve this problem by using PM-SDP to solve for pairwise orthogonal transformations and permutations over the entire dataset and then modifying the ICP procedure we mentioned in section 4 to project onto the set of *consistent orthogonal transformations*; The details of the projection procedure and definition of consistency are given in appendix B. To demonstrate the flexibility of our approach, we use a variation of the high dimensional embedding used above. We embedded the shapes into a seven-dimensional space, the first three coordinates being the euclidian x, y, z coordinates, and the other 4 were eigenfunctions of the LB operator (as was done for isometric matching above). Since the Euclidian coordinates should not mix with the eigenfunction coordinates we constrain R to be block diagonal.

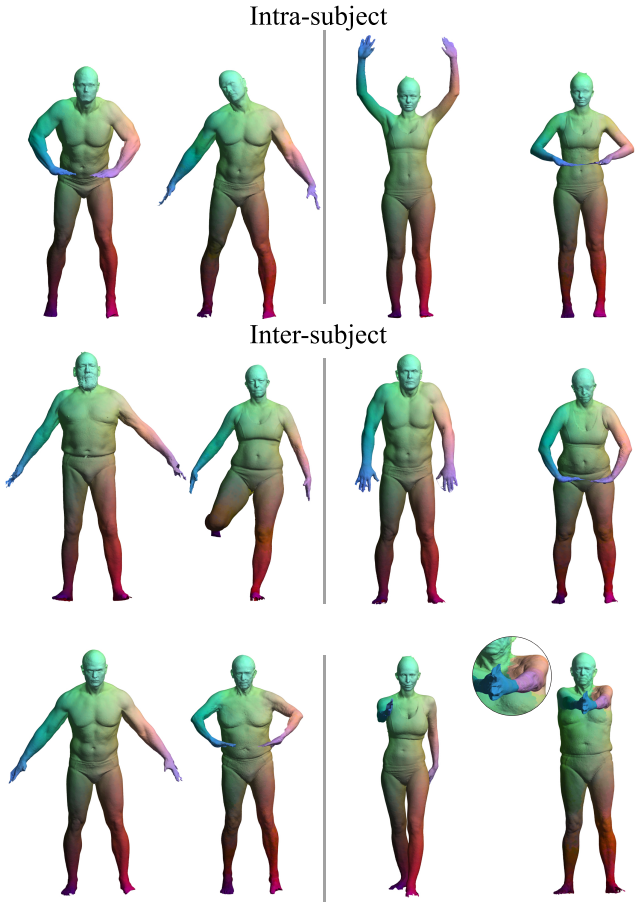

Figure 9: *Examples of typical maps obtained with PM-SDP on the FAUST dataset [Bogo et al., 2014]. Top row: intra-subject. Bottom rows: inter-subject. Bottom right: a failure case.*

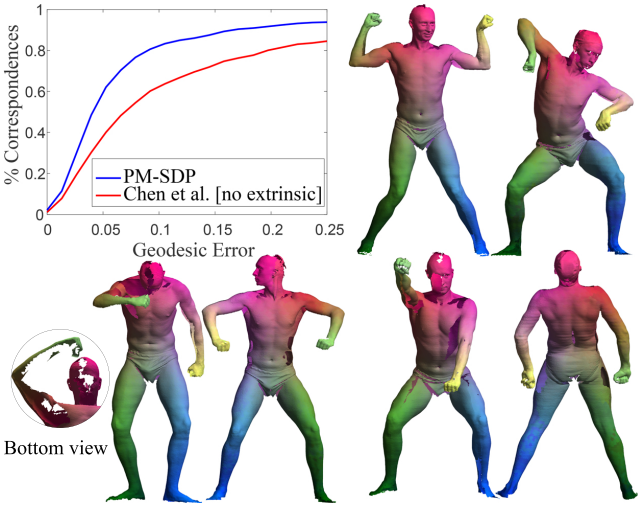

Figure 10: *Performance on the SCAPE raw scans dataset [Anguelov et al., 2005]. Top left: Cumulative error distribution. Other: Examples of typical maps obtained with PM-SDP . Bottom right: a failure case (forward-backward flip).*

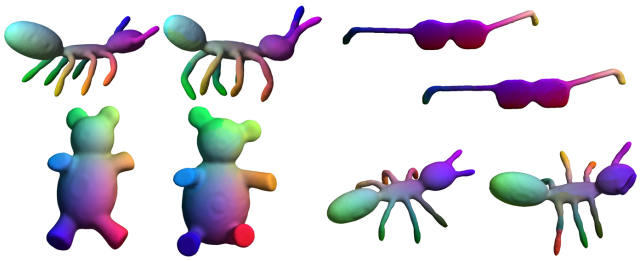

Figure 11: *Examples of maps obtained with PM-SDP on the non-isometric SHREC07 dataset [Giorgi et al., 2007]. Bottom right: a failure case (incorrect corresponding legs).*

As demonstrated in Figure 13 PM-SDP with $d = 7$ (second row) yielded a better consistent alignment in comparison with the method for $d = 3$. The shapes for this experiment are taken from three classes of the SHREC07 [Giorgi et al., 2007] dataset. We made sure the shapes are arbitrarily rotated, sampled $n = k = 20$ farthest points on each shape and solved for all pairwise matchings; for $d = 3$ each pair is computed in 2-3 seconds and for $d = 7$ each pair takes 15-20 seconds.

Timing. Timing of experiments that appear in the paper have already been stated. Here we provide quantitative timing experiments. Figure 14 shows typical run times as a function of dimension or number of points. The experiments were conducted on random and noisy synthetic data. In experiment (a) the dimension d varies from 3 to 20 and we match $k = 50$ points to $n = 100$ points. Experiment (b) compares runtime versus the number of points: in each experiment we match a k point point cloud to a $n = 2k$ point point cloud (up to $k = 50$, $n = 100$) and the dimension is constant $d = 10$. In both cases, R was constrained to be 5-diagonal and we allowed each point to be matched to 30% of the points in the other point cloud based on prior knowledge (in this case these points were selected randomly). (c) Shows comparison of the running time of PM-SDP and the full SDP relaxation discussed in section 3. In this case we use $d = 10, k = n = 5 \ldots 25$. Notably, the full relaxation becomes intractable for more than 17 point, whereas the equivalent PM-SDP formulation solves these problems in just seconds.

7 Conclusions

Summary. We have developed an algorithm that approximates the global minimum of the PM problem with a proven exact recovery property in presence of bilateral symmetries, as well as several other theoretical properties of the algorithm. We demonstrated state of the art results for non-rigid isometric and near-isometric shape matching problems solved using our convex relaxation. We also showed that PM-SDP is useful for anatomical classification of shapes and for aligning shape collections.

Limitations. In contrast to previous SDP relaxations of similar problems, we are able to deal with the registration of around one hundred points. Nonetheless, in comparison with non-SDP based approaches, the main limitation of this algorithm remains its time complexity, which we predict will improve as research on SDP optimization progresses; another limitation of our shape matching framework is the fact that spectral embedding is aimed at near-isometric matching, and is not a good model of the problem for non-isometric shapes.

Future work. One direction we intend to pursue is applying our technique for constructing efficient relaxations for quadratic optimization to different problems other than PM. An interesting theoretical problem which we intend to pursue is proving that PM-SDP (or similar relaxations) give a good approximation of the solution in the general (noisy, far from exact) case, in contrast with our theoretical analysis here which applies only for isometric or near-isometric shapes. Extending to two-way partial matching is also interesting.

Dataset	Classification	PM-SDP	Boyer et al.	Expert
Teeth	Genera	**91.9**	90.9	**91.9**
	Family	**94.3**	92.5	**94.3**
	Above Family	**98.2**	94.8	95.7
Metatarsal	Genera	79.6	79.6	**88.1**
	Family	**93.4**	91.8	**93.4**
	Above Family	**100**	**100**	**100**
Radius	Genera	82.2	**84.4**	77.8
	Family	NA	NA	NA
	Above Family	NA	NA	NA

Table 1: *Classification results (accuracy) achieved by PM-SDP on three anatomical shape data sets compared with [Boyer et al., 2011] and a human expert.*

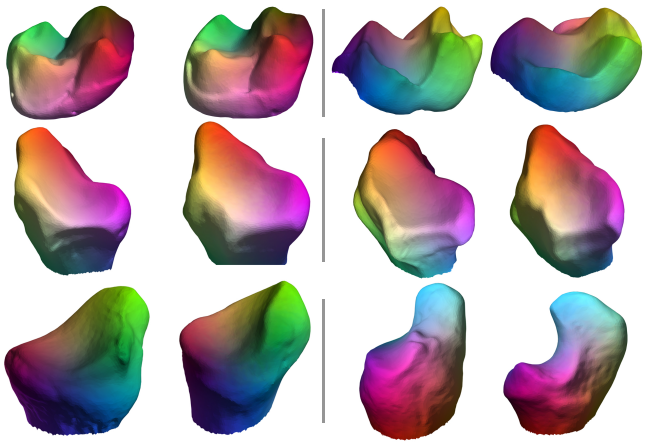

Figure 12: *Examples of typical maps that were obtained by PM-SDP on the anatomical datasets of [Boyer et al., 2011]. First row: Teeth; second row: Metatarsal bone; third row: Radius bone.*

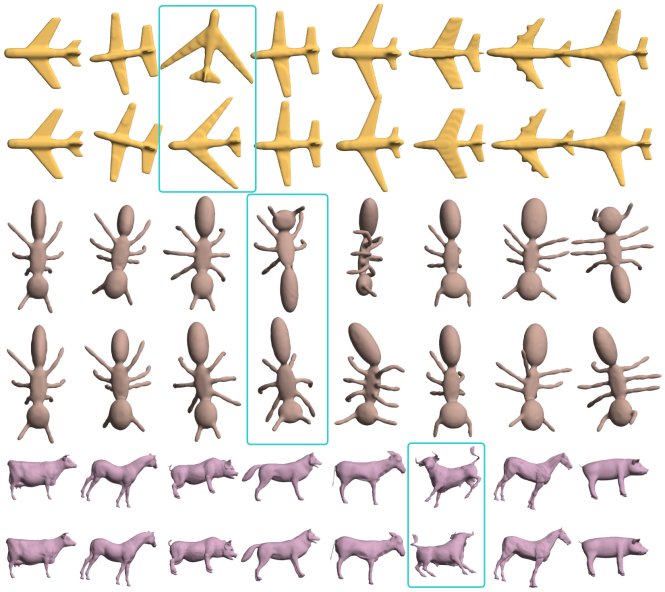

Figure 13: *Finding consistent orthogonal transformations between non-isometric shapes. The figure shows three classes from the SHREC07 dataset [Giorgi et al., 2007]. The first row of each class shows the alignment obtained by PM-SDP in 3D. The second row shows the alignment achieved using a 7-dimensional embedding with PM-SDP. Rectangles show corrected orthogonal transformations.*

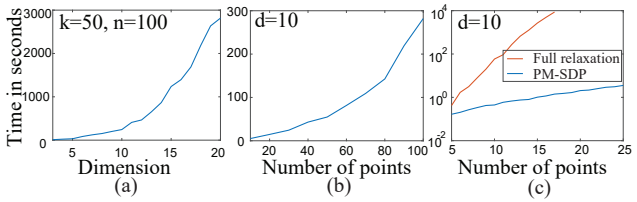

Figure 14: *typical run times as functions of the dimension and the number of points. (a) as a function of the dimension d, (b) as a function of the number of points n, (c) Comparison to the full relaxation discussed in section 3.*

Acknowledgements. This research was supported in part by the European Research Council (ERC Starting Grant "SurfComp", Grant No. 307754), the Israel Science Foundation (Grant No. 1284/12), I-CORE program of the Israel PBC and ISF (Grant No. 4/11). The authors would like thank Maks Ovsjanikov, Vladlen Koltun, Alex and Michael Bronstein and Qifeng Chen for code, helpful discussions and comments, and James Davis for the raw SCAPE scans, and the anonymous reviewers for their useful comments and suggestions.

References

Anguelov, D., Srinivasan, P., Koller, D., Thrun, S., Rodgers, J., and Davis, J. (2005). SCAPE: shape completion and animation of people. In *ACM Transactions on Graphics (TOG)*, volume 24, pages 408–416. ACM.

Aubry, M., Schlickewei, U., and Cremers, D. (2011). The wave kernel signature: A quantum mechanical approach to shape analysis. In *Computer Vision Workshops (ICCV Workshops), 2011 IEEE International Conference on*, pages 1626–1633. IEEE.

Babai, L., Grigoryev, D. Y., and Mount, D. M. (1982). Isomorphism of graphs with bounded eigenvalue multiplicity. In *Proceedings of the Fourteenth Annual ACM Symposium on Theory of Computing*, STOC '82, pages 310–324, New York, NY, USA. ACM.

Belkin, M. and Niyogi, P. (2003). Laplacian eigenmaps for dimensionality reduction and data representation. *Neural computation*, 15(6):1373–1396.

Berg, A. C., Berg, T. L., and Malik, J. (2005). 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, pages 26–33. IEEE.

Besl, P. and McKay, N. (1992). A method for registration of 3-d shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2):239–256.

Bogo, F., Romero, J., Loper, M., and Black, M. J. (2014). FAUST: Dataset and evaluation for 3D mesh registration. In *Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on*, pages 3794–3801. IEEE.

Boyer, D. M., Lipman, Y., Clair, E. S., Puente, J., Patel, B. A., Funkhouser, T., Jernvall, J., and Daubechies, I. (2011). Al-

gorithms to automatically quantify the geometric similarity of anatomical surfaces. *Proceedings of the National Academy of Sciences*, 108(45):18221–18226.

Bronstein, A. M., Bronstein, M. M., and Kimmel, R. (2006). 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):1168–1172.

Brown, B. J. and Rusinkiewicz, S. (2007). Global non-rigid alignment of 3-d scans. *ACM Transactions on Graphics (TOG)*, 26(3):21.

Chen, Q. and Koltun, V. (2015). Robust nonrigid registration by convex optimization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2039–2047.

Dym, N. and Lipman, Y. (2016). Exact recovery with symmetries for procrustes matching. *Technical report*.

Eldar, Y., Lindenbaum, M., Porat, M., and Zeevi, Y. Y. (1997). The farthest point strategy for progressive image sampling. *Image Processing, IEEE Transactions on*, 6(9):1305–1315.

Fischler, M. A. and Bolles, R. C. (1981). Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395.

Fukuda, M., Kojima, M., Murota, K., and Nakata, K. (2000). Exploiting sparsity in semidefinite programming via matrix completion I: General framework. *SIAM J. on Optimization*, 11(3):647–674.

Gelfand, N., Mitra, N. J., Guibas, L. J., and Pottmann, H. (2005). Robust global registration. In *Proceedings of the Third Eurographics Symposium on Geometry Processing*, SGP '05, Aire-la-Ville, Switzerland, Switzerland. Eurographics Association.

Giorgi, D., Biasotti, S., and Paraboschi, L. (2007). Shape retrieval contest 2007: Watertight models track. *SHREC competition*, 8.

Gower, J. C. and Dijksterhuis, G. B. (2004). *Procrustes problems*, volume 3. Oxford University Press Oxford.

Grone, R., Johnson, C. R., Sá, E. M., and Wolkowicz, H. (1984). Positive definite completions of partial Hermitian matrices. *Linear algebra and its applications*, 58:109–124.

Huang, Q., Wang, F., and Guibas, L. (2014). Functional map networks for analyzing and exploring large shape collections. *ACM Trans. Graph.*, 33(4):36:1–36:11.

Jain, V., Zhang, H., and Van Kaick, O. (2007). Non-Rigid Spectral Correspondence of Triangle Meshes. *International Journal of Shape Modeling*, 13:101–124.

Kendall, D. G. (1984). Shape manifolds, Procrustean metrics, and complex projective spaces. *Bulletin of the London Mathematical Society*, 16(2):81–121.

Kezurer, I., Kovalsky, S. Z., Basri, R., and Lipman, Y. (2015). Tight relaxation of quadratic matching. *Computer Graphics Forum*, 34(5):115–128.

Kim, V. G., Lipman, Y., and Funkhouser, T. (2011). Blended intrinsic maps. *ACM Transactions on Graphics*, 30(4):1.

Leordeanu, M. and Hebert, M. (2005). A spectral technique for correspondence problems using pairwise constraints. In *Computer Vision, 2005. ICCV 2005. Tenth IEEE International Conference on*, volume 2, pages 1482–1489. IEEE.

Li, H., Sumner, R. W., and Pauly, M. (2008). Global correspondence optimization for non-rigid registration of depth scans. In *Computer Graphics Forum*, volume 27, pages 1421–1430. Wiley Online Library.

Lipman, Y. and Funkhouser, T. (2009). Möbius voting for surface correspondence. In *ACM Transactions on Graphics (TOG)*, volume 28, page 72. ACM.

Luo, Z.-Q., Ma, W.-k., So, A. M.-C., Ye, Y., and Zhang, S. (2010). Semidefinite relaxation of quadratic optimization problems. *Signal Processing Magazine, IEEE*, 27(3):20–34.

Mémoli, F. and Sapiro, G. (2004). Comparing point clouds. In *Proceedings of the 2004 Eurographics/ACM SIGGRAPH Symposium on Geometry Processing*, SGP '04, pages 32–40, New York, NY, USA. ACM.

Mitteroecker, P. and Gunz, P. (2009). Advances in geometric morphometrics. *Evolutionary Biology*, 36(2):235–247.

MOSEK (2015). *The MOSEK optimization toolbox for MATLAB manual. Version 7.1 (Revision 49).*

Nguyen, A., Ben-Chen, M., Welnicka, K., Ye, Y., and Guibas, L. (2011). An optimization approach to improving collections of shape maps. *Computer Graphics Forum*, 30(5):1481–1491.

Ovsjanikov, M., Ben-Chen, M., Solomon, J., Butscher, A., and Guibas, L. (2012). Functional maps: a flexible representation of maps between shapes. *ACM Transactions on Graphics (TOG)*, 31(4):30.

Ovsjanikov, M., Mérigot, Q., Mémoli, F., and Guibas, L. (2010). One point isometric matching with the heat kernel. In *Computer Graphics Forum*, volume 29, pages 1555–1564. Wiley Online Library.

Ovsjanikov, M., Sun, J., and Guibas, L. (2008). Global intrinsic symmetries of shapes. In *Computer Graphics Forum*, volume 27, pages 1341–1348. Wiley Online Library.

Pinkall, U. and Polthier, K. (1993). Computing discrete minimal surfaces and their conjugates. *Experimental mathematics*, 2(1):15–36.

Pokrass, J., Bronstein, A. M., Bronstein, M. M., Sprechmann, P., and Sapiro, G. (2013). Sparse modeling of intrinsic correspondences. *Computer Graphics Forum*, 32(2 PART4):459–468.

Poljak, S., Rendl, F., and Wolkowicz, H. (1995). A recipe for semidefinite relaxation for (0, 1)-quadratic programming. *Journal of Global Optimization*, 7(1):51–73.

Rusinkiewicz, S. and Levoy, M. (2001). Efficient variants of the ICP algorithm. *Proceedings Third International Conference on 3-D Digital Imaging and Modeling*, pages 145–152.

Rustamov, R. M. (2007). Laplace-Beltrami eigenfunctions for deformation invariant shape representation. In *Proceedings of the Fifth Eurographics Symposium on Geometry Processing*, SGP '07, pages 225–233, Aire-la-Ville, Switzerland, Switzerland. Eurographics Association.

Saunderson, J., Parrilo, P. A., and Willsky, A. S. (2014). Semidefinite descriptions of the convex hull of rotation matrices. *arXiv preprint arXiv:1403.4914*.

Singer, A. and Wu, H.-T. (2012). Vector diffusion maps and the connection laplacian. *Communications on Pure and Applied Mathematics*, 65(8):1067–1144.

Tam, G. K., Cheng, Z.-Q., Lai, Y.-K., Langbein, F. C., Liu, Y., Marshall, D., Martin, R. R., Sun, X.-F., and Rosin, P. L. (2013). Registration of 3D point clouds and meshes: a survey from rigid to nonrigid. *Visualization and Computer Graphics, IEEE Transactions on*, 19(7):1199–1217.

Tevs, A., Bokeloh, M., Wand, M., Schilling, A., and Seidel, H.-P. (2009). Isometric registration of ambiguous and partial data. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 1185–1192. IEEE.

Van Kaick, O., Zhang, H., Hamarneh, G., and Cohen-Or, D. (2011). A survey on shape correspondence. In *Computer Graphics Forum*, volume 30, pages 1681–1707. Wiley Online Library.

Waki, H., Kim, S., Kojima, M., and Muramatsu, M. (2006). Sums of squares and semidefinite programming relaxations for polynomial optimization problems with structured sparsity. *SIAM Journal on Optimization*, 17:218–242.

Wand, M., Chang, W., Li, H., Mitra, N., Pauly, M., and Rusinkiewicz, S. (2011). Computing correspondences in geometric data sets tutorial. http://resources.mpi-inf.mpg.de/deformableShapeMatching/EG2011_Tutorial/.

Wei, L., Huang, Q., Ceylan, D., Vouga, E., and Li, H. (2015). Dense Human Body Correspondences Using Convolutional Networks.

Yang, J., Li, H., and Jia, Y. (2013). Go-ICP: Solving 3D Registration Efficiently and Globally Optimally. *2013 IEEE International Conference on Computer Vision*, pages 1457–1464.

Zeng, Y., Wang, C., Wang, Y., Gu, X., Samaras, D., and Paragios, N. (2010). Dense non-rigid surface registration using high-order graph matching. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, pages 382–389. IEEE.

Zhao, Q., Karisch, S. E., Rendl, F., and Wolkowicz, H. (1998). Semidefinite programming relaxations for the quadratic assignment problem. *Journal of Combinatorial Optimization*, 2(1):71–109.

Zuffi, S. and Black, M. J. (2015). The stitched puppet: A graphical model of 3D human shape and pose. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3537–3546.

Appendix A Local minimization

The local minimization can be initialized using the output of PM-SDP in four different ways. Two immediate possibilities are given by choosing the R or X coordinates from the optimal solution. We note that R, X may not be in $\mathcal{O}(d) \times \Pi_n$, but as mentioned previously they are in the convex hull of that set. Two other possibilities can be obtained from decomposing the lifted variables B_j as explained next. When (7) holds, we can use (8) to write $B_j = [R]\, X_j^T$. We can then combine the matrices B_j into a larger matrix B so that $B = [R]\,[X]^T$, and X, R can be recovered in this case by factorizing B into an outer product of two vectors. Motivated by this, we do the following:

1. Create the matrix B using the matrices B_j: $B = [B_1, ..., B_k]$.
2. Project B onto the set of rank one matrices using SVD, denote the projection as $\widehat{B}$.
3. Factorize the projected rank-one matrix as an outer product of two vectors $\widehat{B} = [\widehat{R}][\widehat{X}]$, and use these vectors as another two possible initializations of R and X.

Appendix B Collection alignment

In this section we present the collection alignment algorithm mentioned in sub-section 6.3 . First we solve PM for each pair of shapes in the collection with the additional constraint that R_{ij} is in the convex hull of $SO(3)$ using the formula from [Saunderson et al., 2014]. Then, we build a large matrix R containing all the orthogonal transformations obtained by PM-SDP as sub-blocks. More specifically, each $d \times d$ block in the (i, j)-th position is the orthogonal transformation between shape i and shape j. In case the set of transformations is consistent (i.e., for each i, j, k, $R_{ij} \cdot R_{jk} = R_{ik}$) , R is positive semidefinite and has rank d (for the definition of consistent maps and orthogonal transformations see [Singer and Wu, 2012; Huang et al., 2014; Nguyen et al., 2011]).

In the spirit of this observation, we feed this matrix into a an iterative ICP-like algorithm that performs the following steps:

1. Project R onto the set of consistent rank d matrices: Let UDU^T be the eigen-decomposition of R, we take the largest d eigenvectors of U scaled by the square root of the corresponding eigenvalue, which we denote U_d and project each of its $d \times d$ block to its closest orthogonal matrix. Denote the matrix with the new blocks as U'_d then the output of this step is $R' = U'_d(U'_d)^T$.
2. For each block of R' solve for the best permutation X_{ij} using linear programming as described in section 4 .
3. For each permutation X_{ij} solve for the best orthogonal transformation using the closed form solution mentioned above.
4. Construct R from the matrices R_{ij}.
5. Iterate until convergence.

Our experiments show that this algorithm reaches a steady state after a few iterations that take a few seconds for the problems in Subsection 6.3.