

# Simplicial Complex based Point Correspondence between Images warped onto Manifolds

Charu Sharma and Manohar Kaul

Indian Institute of Technology Hyderabad, India  
{cs16resch11007,mkaul}@iith.ac.in

**Abstract.** Recent increase in the availability of warped images projected onto a curved manifold, especially omnidirectional spherical ones, coupled with the success of higher-order assignment methods, has sparked an interest in the search for improved higher-order matching algorithms on warped images due to projection. Although, currently, several existing methods “flatten” such 3D images to use planar graph / hypergraph matching methods, they still suffer from severe distortions and other undesired artifacts, which result in inaccurate matching. Alternatively, current planar methods cannot be trivially extended to effectively match points on images warped on curved manifold. Hence, matching on these warped images persists as a formidable challenge. In this paper, we pose the assignment problem as finding a bijective map between two graph induced simplicial complexes, which are higher-order analogues of graphs. We propose a constrained quadratic assignment problem (QAP) that matches each  $p$ -skeleton of the simplicial complexes, iterating from the highest to the lowest dimension. The accuracy and robustness of our approach are illustrated on both synthetic and real-world spherical / warped (projected) images with known ground-truth correspondences. We significantly outperform existing state-of-the-art spherical matching methods on a diverse set of datasets.

## 1 Introduction

There exists a longstanding line of research on finding bijective correspondences (i.e., assignments / matchings<sup>1</sup>) between two sets of visual features. Notable applications include stereo matching [14], structure from motion (SfM) [34], and image registration [31], to name a few. Traditionally, when matching points between multiple images of a fixed environment from various viewpoints, most approaches recover matchings and relative camera geometry (e.g. fundamental matrix) using a robust technique such as RANSAC [16]. On the other hand, when matching between different instances of the same category, graph matching methods [38] using *unary* and *pairwise* constraints have been successfully utilized. More recently, graph matching has been subsumed by *hypergraph matching* using *higher-order* constraints [12,19]. An important appeal of higher-order matching methods is their ability to coherently match compact local geometric

---

<sup>1</sup> *assignment* and *matching* are used interchangeably in this paper.

features from the source space to similar compact regions in the target space, despite the presence of noise, outliers, and incomplete data, thus achieving accurate matches that are also *local structure-preserving* in nature.

The recent proliferation of spherical images (e.g., omnidirectional and panoramic images captured from cameras mounted on drones and autonomous vehicles) and more generally, images warped onto *curved manifolds*, has sparked a heightened interest in assignment algorithms on such datasets due to the challenges they present in terms of curvature, both uniform and non-uniform [32,18,36,35]. Although assignment problems have been well studied for decades in computer vision, a majority of the work has only focused on matching points between *planar (flat) images*. Therefore, matching points on images with warping transformations which fall into the category of projective parametric models remains a challenging task, mainly due to the introduction of undesirable artifacts like severe distortions in pairwise distances between landmark points, non-linear distortions in local geometries, noise, illumination, blur, and occlusions [3,8], when flattening.

When dealing with matchings on curved geometries, primarily two types of methods are employed. Some putative matchings are computed to estimate a *fundamental matrix* [16,10] that captures the *epipolar geometry* of the 3D image. *Stereo rectification* [5] uses this fundamental matrix to re-project the two images on the same flat plane with row images aligned in parallel, followed by a re-matching to improve matching accuracy. Alternatively, *geometric alignment* on the fundamental matrix is used to *verify* and distinguish *inliers* from *outliers*, so that outliers can be pruned post matching to further boost accuracy [34]. Elements warped on the curved manifold cannot be metrically sampled in such methods and hence severe distortions are introduced [7], which is also consistent with the findings in our empirical studies.

**Applications** An interesting and noteworthy application of higher-order matching on spherical images arises in the area of biomedical imaging, especially in *retinal imaging* using *optical coherence tomography* (OCT). To investigate a wider *field of view*, 3D *fundus* images of the eye are captured, matched, and “stitched” together to form an *OCT montage* [22,25]. This matching operation must additionally preserve *regions of interest* such as the *optic cup / disc*, *fovea*, *macula*, *vessels*, and *microaneurysms*, to name a few [29]. In addition to the standard noise, occlusion, and artifacts in these OCT fundus images, the data also suffers from data shifts due to axial eye motions and unpredictability between *eye positions* and *instrument alignment* across various scans [22]. Therefore, OCT datasets cannot easily be matched using rigid 3D transformations. Such images are not limited to merely spherical ones, but also arise in more general warped images due to projection. For instance, 3D sonograms depict the cervix as a *conic frustum (truncated cone)* [1] and clustered nanofluid microflow patterns in elastic micro-tubes are tracked via matching between *cylindrical* images in a time-lapse [33].

**Our method** In this paper, we focus on exploiting the intrinsic higher-order geometric relationships between landmark points on images warped onto curved

manifolds. We capture these higher-order connections by constructing a combinatorial topological structure (simplicial complex) which is induced by a graph, whose *vertices* are the landmark points embedded on the warped image and whose *edges* are *geodesic curves* between selected vertex pairs. Next, we pose the assignment problem as a multi-dimensional *quadratic assignment problem (QAP)* between two graph-induced simplicial complexes.

**Our contributions** (i) To the best of our knowledge, we are the first to propose matching landmark points on warped images projected onto curved manifolds. (ii) In an attempt to break away from other works which solely focus on flat or spherical images, we propose a novel graph induced simplicial complex that efficiently captures higher order structures in a succinct manner, considering the inherent properties of the underlying curved manifold on which the landmark points are embedded. (iii) We uniquely formulate the assignment problem as a multi-dimensional combinatorial matching between two graph induced simplicial complexes, propose a novel algorithm to solve it, and analyze the time-complexity of our algorithm. (iv) Finally, to illustrate the robustness of our proposed method, we perform extensive experiments by comparing to planar matching methods, both *existing* and *extended by us* as *naive baselines* for matching on manifolds. We compare our method against existing *graph matching* and *spherical matching* (both boosted using *rectification* and *verification* techniques) [21,27,37,38,39,30,12] on warped images and interestingly observe that not only does our method significantly outperform these matching methods on warped images onto curved manifolds (with up to 49.7% matching error reduction), but it also outperforms existing planar matching algorithms on “flat” planar images too (with up to 42.2% matching error reduction), due to the ability to naturally capture higher-order relationships by the simplicial complex.

## 2 Preliminaries

In this section, we introduce our notation and provide the necessary background for our higher-order assignment algorithm on curved manifolds. We begin by introducing certain standard definitions followed by our problem definition.

Let  $\mathcal{M}$  denote a *curved manifold*, i.e., with zero genus and no boundaries. On a plane, the shortest distance between any two points is a straight line, i.e., a curve whose derivative to its tangent vectors is zero. We extend this notion of a “straight line” to *curved manifolds* by defining the shortest path (on  $\mathcal{M}$ ) between its endpoints  $u$  and  $v$  placed on  $\mathcal{M}$ , as a *geodesic curve*  $\gamma(u, v)$ .

**Simplicial complex** We begin by providing some general definitions before we can formally define a simplicial complex. More background can be found in [24].

Given a set  $V = \{v_0, \dots, v_n\}$  of  $(n+1)$  affinely independent points in  $\mathbb{R}^{n+1}$ , a *n-dimensional simplex* (also called *n-simplex*)  $\sigma^{(n)}$  with *vertices*  $V$  is the *convex hull* of  $V$ , i.e., more formally

$$\sigma^{(n)} = \left\{ (k_0, k_1, \dots, k_n) \in \mathbb{R}^{n+1} \mid \sum_{i=0}^n k_i = 1, k_i \geq 0 \forall i \right\}$$

The dimension of  $n$ -simplex  $\sigma^{(n)}$  is denoted by  $\dim(\sigma^{(n)})$ . For example, a *point* / *vertex* (0-simplex), an *edge* (1-simplex), and a *triangle* (2-simplex) are represented as  $\sigma^{(0)}$ ,  $\sigma^{(1)}$ , and  $\sigma^{(2)}$ , respectively. For  $0 \leq i \leq n$ , the  $i$ -th *facet*  $f_i$  of the  $n$ -simplex  $\sigma^{(n)}$  is the  $(n-1)$ -simplex  $\sigma^{(n-1)}$ , whose vertices are those underlying  $\sigma^{(n)}$ , except the  $i$ -th vertex. For example, a 2-simplex (triangle) has three 1-simplices (edges) as *facets*. The *boundary*  $\partial\sigma^{(n)}$  of the  $n$ -simplex  $\sigma^{(n)}$  is  $\bigcup_{i=0}^n f_i$ . Finally, a *simplicial complex*  $\mathcal{K}$  is a set of simplices that satisfy the following conditions: (i). Any face of a simplex in  $\mathcal{K}$  is a simplex in  $\mathcal{K}$  and (ii). Intersection of distinct simplices  $\sigma_i$  and  $\sigma_j$  in  $\mathcal{K}$ , is a *common face* of both  $\sigma_i$  and  $\sigma_j$ <sup>2</sup>. The  $p$ -*skeleton*  $\mathcal{K}^{(p)} \subset \mathcal{K}$  is formed by the set of  $k$ -simplices  $\sigma^{(k)}$ , where  $k \leq p$ . Additionally, we denote by  $\mathcal{K}_k$  the set of  $k$ -simplices in  $\mathcal{K}$ . The dimension  $\dim(\mathcal{K})$  of a simplicial complex  $\mathcal{K}$  is the maximum of the dimensions of its constituent simplices.

**Problem definition** Our problem consists of first constructing *geometric simplicial complexes* between landmark points given on curved manifolds, followed by finding an optimal (i.e., least cost) assignment between a pair of such geometric simplicial complexes by matching simplices of the same dimension, one dimension at a time. More formally, Let  $P$  and  $P'$  denote two sets of *landmark points* on curved manifolds  $\mathcal{M}$  and  $\mathcal{M}'$ , respectively. We construct *geometric simplicial complexes*  $\mathcal{K}$  and  $\mathcal{K}'$  whose set of vertices (0-simplices) are  $P$  and  $P'$ . The edges/arcs (1-simplices) in  $\mathcal{K}$  and  $\mathcal{K}'$  are given by geodesics between select few pairs of vertices, from their corresponding vertex sets.

Given two simplicial complexes  $\mathcal{K}$  and  $\mathcal{K}'$ , we assume without loss of generality, that the number of simplices of each corresponding dimension are equal in both complexes. Then, our goal is to find a set of  $h$  bijective *matching functions*  $\{m_k\}_{k=0}^h : \mathcal{K} \rightarrow \mathcal{K}'$  that match the set of  $k$ -simplices in  $\mathcal{K}$  (i.e.,  $\mathcal{K}_k$ ) to  $k$ -simplices in  $\mathcal{K}'$  (i.e.,  $\mathcal{K}'_k$ ), for dimensions  $k = 0 \dots h$ , to minimize the overall objective function

$$\operatorname{argmin}_{m_0, \dots, m_h} \sum_{k=0}^h \sum_{i=1}^{|\mathcal{K}_k|} c(\sigma_i^{(k)}, m_k(\sigma_i^{(k)})) \quad (1)$$

where  $c(\cdot, \cdot)$  is the *geometric matching cost* between a  $k$ -simplex  $\sigma^{(k)}$  in  $\mathcal{K}$  to a  $k$ -simplex  $m(\sigma^{(k)})$  in  $\mathcal{K}'$  and simplicial complex dimension  $h = \min(\dim(\mathcal{K}), \dim(\mathcal{K}'))$ . Unlike formulations proposed in graph matching methods [38], where only node and pairwise geometric relations are considered, our *combinatorial optimization* formulation takes into consideration higher-order geometric constraints, which better excludes ambiguous matchings. In subsequent sections, we show how we construct such geometric simplicial complexes from the landmark points on curved manifolds (Section 3), followed by a detailed explanation of our assignment algorithm (Section 4).

<sup>2</sup> For ease of notation, we drop the dimension superscript and index subscript for a simplex when it is understood from context.

### 3 Building a Simplicial Complex on a Curved Manifold

In this section, inspired by the work of Dey et. al. [11], we similarly construct a *graph-induced* simplicial complex, which is built upon a graph connecting the landmark points. We begin by describing the process of constructing the *underlying graph*.

**Graph construction** Let  $(P, g)$  denote the set of landmark points  $P$  with a metric  $g$  that denotes the geodesic distance between a pair of points on  $\mathcal{M}$ . Additionally, let the  $k$ -neighborhood  $\mathcal{N}_k(u)$  denote the set of  $k$  nearest neighbors of landmark point  $u \in P$  (inclusive of  $u$ ) on manifold  $\mathcal{M}$  according to the geodesic metric  $g$ .

Considering all ordered pairs  $(u, v)$ , where  $u, v \in P$ , an undirected *edge/arc* is introduced between points  $u$  and  $v$ , when their corresponding  $k$ -neighborhoods  $\mathcal{N}_k(u)$  and  $\mathcal{N}_k(v)$  have a non-empty intersection, i.e.,  $\mathcal{N}_k(u) \cap \mathcal{N}_k(v) \neq \emptyset$ . All such edges are collected into a set denoted by  $E$ . This completes the construction of our underlying graph  $G = (P, E)$ . Observe that the vertex set (landmarks)  $P$  form the 0-skeleton  $\mathcal{K}^{(0)}(G)$  and the sets  $E$  and  $P$  together form the 1-skeleton  $\mathcal{K}^{(1)}(G)$ , of our graph-induced simplicial complex that we will denote by  $\mathcal{K}(G)$ .

Recall that a  $n$ -clique in a graph is a complete subgraph between  $n$  vertices, i.e., it consists of  $n$  vertices and  $\binom{n}{2}$  edges.

**Graph-induced complex**  $\mathcal{K}(G)$  is defined as the simplicial complex where a  $n$ -simplex  $\sigma^{(n)} = \{p_1, p_2, \dots, p_{n+1}\}$  is in  $\mathcal{K}(G)$ , if and only if there exists a  $(n+1)$ -clique  $\{p_1, p_2, \dots, p_{n+1}\} \subseteq P$  in the underlying graph  $G = (P, E)$ . In words, the *cliques* of the underlying graph  $G = (P, E)$  form the *simplices* in  $\mathcal{K}(G)$  because cliques satisfy both conditions of being a simplicial complex (which can be trivially verified). In order to be used in our assignment algorithm, we must represent the graph-induced simplicial complex  $\mathcal{K}(G)$  as a set of *boundary matrices*, which we present next.

**Matrix representation of  $\mathcal{K}(G)$ :** Given  $\mathcal{K}(G)$  and its  $p$ -skeleton  $\mathcal{K}^{(p)}(G)$  that contains cliques upto size  $p+1$ , we represent it as a *boundary matrix*  $M_p \in \mathbb{Z}^{n \times m}$  defined as

$$M_p = \begin{matrix} & \sigma_1^{(p)} & \dots & \sigma_m^{(p)} \\ \tau_1^{(p-1)} & \left( \begin{matrix} a_{11} & \dots & a_{1m} \\ \vdots & \ddots & \vdots \\ a_{n1} & \dots & a_{nm} \end{matrix} \right) \\ \vdots & & & \\ \tau_n^{(p-1)} & & & \end{matrix}$$

where  $a_{ij} = 1$  if and only if the  $i$ -th  $(p-1)$ -simplex  $\tau_i^{(p-1)}$  is a *facet* of the  $j$ -th  $p$ -simplex  $\sigma_j^{(p)}$ , otherwise  $a_{ij} = 0$ . Then, the boundary of a  $j$ -th  $p$ -simplex is given by  $\partial_p \sigma_j^{(p)} = \sum_{i=1}^n a_{ij} \tau_i^{(p-1)}$ .

Observe that the  $p$ -th boundary matrix  $M_p$  captures all possible relationships between  $p$ -simplices and their  $(p-1)$ -simplex boundaries (or facets). Boundary matrix  $M_p$  is made for each  $p$ -skeleton and therefore  $\mathcal{K}(G)$  is expressed as a set of boundary matrices  $\{M_p\}_{p=1}^h$ , where  $h = \dim(\mathcal{K}(G))$ .

*Remark 1.* Our underlying graph  $G$  already contains as a *subgraph* a simple  $k$ -nearest neighbor graph which is constructed by introducing edges between a

---

**Algorithm 1** Matching graph induced simplicial complexes

---

**Input:**  $\mathcal{K}(G) = \{M_p\}_{p=1}^h$  and  $\mathcal{K}(G') = \{M'_p\}_{p=1}^h$ 

- 1: **for**  $p = h \dots 1$  **do**
- 2:   Build cost matrix  $\mathcal{L}^{(p)}$  for  $M_p$  and  $M'_p$  (\*account for  $\mathcal{L}^{(p-1)}$ )
- 3:    $X_p^* := \mathbf{Solve\ QAP}$  ( $M_p, M'_p, \mathcal{L}^{(p)}$ )
- 4:    $\mathcal{L}^{(p-1)} :=$  Build cost matrix of  $(p-1)$ -faces
- 5:   from successful  $p$ -simplex matches.
- 6: **end for**

**Return:**  $\{X_1^*, \dots, X_h^*\}$  # set of permutation matrices

---

vertex in question and its  $k$  nearest neighbors. Therefore, our underlying graph  $G$  has more edges and thus has a higher likelihood to form higher-order relations between vertices. On the other hand, while the Delaunay triangulation is simple to compute and is a good vehicle for extracting topology of sampled spaces, its size becomes prohibitively large for reasonable computations and thus adversely affects the QAP matching algorithm.

In summary, our underlying graph  $G$  which is inspired by the *Vietoris-Rips* complex construction provides a good *proximity structure*, which is neither *too sparse* (like simple  $k$ -NN graphs) or *too dense* (like Delaunay triangulated graphs) and encodes useful higher-order information about local relations of points in  $P$ .

## 4 Assignment Algorithm

Recall our problem definition (Section 2) of trying to find a set of assignments / matching functions between two graph-induced simplicial complexes  $\mathcal{K}(G)$  and  $\mathcal{K}(G')$ . Here, we outline the details of our assignment algorithm.

Given a boundary matrix  $M_p \in \mathbb{Z}^{n \times m}$  that represents a  $p$ -skeleton  $\mathcal{K}^{(p)}(G)$ , we first capture the *geodesic neighborhood geometry* of simplices in  $M_p$ . We begin by defining an *adjacency operator*  $\sim$  between two simplices followed by a definition of a *neighborhood of a simplex*. This neighborhood of a simplex is then elegantly captured by *affine weight vectors*, which are later used in the matching algorithm.

**Definition 1 (adjacency relation).** *Given two simplices  $\sigma^{(d)}$  and  $\sigma'^{(d')}$ , each of arbitrary dimension  $d$  and  $d'$ , we consider them to be adjacent to one another if and only if they share a common simplex. We denote this adjacency relation by  $\sigma^{(d)} \sim \sigma'^{(d')}$ . The dimension of the common simplex can take values from 0 to  $\min(d, d')$ .*

For example, two 2-simplices / triangles  $\sigma^{(2)}$  and  $\sigma'^{(2)}$  could either be connected at a common 0-simplex / vertex or share a common 1-simplex / edge; both cases would result in the simplices being *adjacent*, i.e.,  $\sigma^{(2)} \sim \sigma'^{(2)}$ .

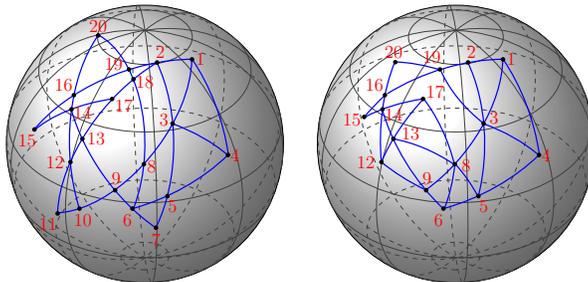
**Simplex neighborhood** The boundary matrix  $M_p$ 's columns encode  $p$ -simplices  $\sigma_1^{(p)}, \dots, \sigma_m^{(p)}$  and its rows encode  $(p-1)$ -simplices  $\tau_1^{(p-1)}, \dots, \tau_n^{(p-1)}$ . The computation of the neighborhood  $\mathfrak{N}(\cdot)$  for  $p$ -simplices and  $(p-1)$ -simplices differ slightly. The neighborhood of a  $p$ -simplex consists of  $p$ -simplices (same dimension) and  $(p-1)$ -simplices (one dimension lower) that are adjacent to it. While, the neighborhood of a  $(p-1)$ -simplex consists of  $(p-1)$ -simplices (same dimension) and  $p$ -simplices (one dimension higher) that are adjacent to it. More formally, the neighborhood of the  $i$ -th  $p$ -simplex  $\sigma_i^{(p)}$  is  $\mathfrak{N}(\sigma_i^{(p)}) = \{\sigma_j^{(p)} \mid \sigma_j^{(p)} \sim \sigma_i^{(p)}\} \cup \{\tau_j^{(p-1)} \mid \tau_j^{(p-1)} \sim \sigma_i^{(p)}\}$  and the neighborhood of the  $i$ -th  $(p-1)$ -simplex  $\tau_i^{(p-1)}$  is  $\mathfrak{N}(\tau_i^{(p-1)}) = \{\tau_j^{(p-1)} \mid \tau_j^{(p-1)} \sim \tau_i^{(p-1)}\} \cup \{\sigma_j^{(p)} \mid \sigma_j^{(p)} \sim \tau_i^{(p-1)}\}$ . Such neighborhoods are computed for all the  $p$ - and  $(p-1)$ -simplices in  $M_p$ , where  $i \neq j$ .

**Affine weight vectors** For a  $p$ -simplex  $\sigma^{(p)}$ , let  $\mathcal{B}(\sigma^{(p)})$  denote the set of all the *barycenters*  $\{b_1, \dots, b_{|\mathfrak{N}(\sigma^{(p)})|}\}$  of the simplices in the neighborhood  $\mathfrak{N}(\sigma^{(p)})$ . Then,  $\sigma^{(p)}$  is represented as an *affine combination* of the barycenters in  $\mathcal{B}(\sigma^{(p)})$ , i.e.,  $\sum_{i=1}^{|\mathcal{B}(\sigma^{(p)})|} \alpha_i b_i$ , where  $\sum_{i=1}^{|\mathcal{B}(\sigma^{(p)})|} \alpha_i = 1$  (i.e., weights  $\alpha_i$ 's must sum to 1). Therefore,  $\sigma^{(p)}$  is expressed as an *affine weight vector*  $\alpha(\sigma^{(p)})$  of dimension  $(n+m)$ , with  $|\mathcal{B}(\sigma^{(p)})|$  positions corresponding to  $\mathfrak{N}(\sigma^{(p)})$  filled with non-empty affine weights and the rest set to zero. Such an affine weight vector is computed for every simplex of dimension  $p$  and  $(p-1)$  contained in  $M_p$ . Among all possible affine representations of a simplex, we chose to use *least squares* to guarantee minimal error under L2-norm, and furthermore it assigns non-zero weights to each of its adjacent simplex barycenters, thereby better capturing the local geometric properties in its neighborhood.

*Remark 2.* The affine weight vectors act as *locally affine invariant descriptors* that can handle complex and natural transformations of the underlying manifold  $\mathcal{M}$ . Additionally, it allows for much fewer variables and can be much more easily linearized in the subsequent QAP formulation. Furthermore, the inclusion of barycenters from neighborhoods of each simplex act as higher-order geometric constraints that easily excludes ambiguous matchings. In comparison, simple matching models that rely on just a distance matrix with pairwise geodesic distances on the manifold are not invariant to local and global affine transformations and completely disregard higher-order relationships.

**Cost matrix construction** Next, we describe the construction of a *cost matrix* that is needed to compute assignments between  $M_p \in \mathbb{Z}^{n \times m}$  and  $M'_p \in \mathbb{Z}^{n' \times m'}$ . We begin by constructing two cost matrices  $\mathcal{C}^{(p-1)} \in \mathbb{R}^{n \times n'}$  and  $\mathcal{C}^{(p)} \in \mathbb{R}^{m \times m'}$  to measure the Euclidean distance between the affine weight vectors of  $(p-1)$ -simplices and the Euclidean distance between the affine weight vectors of  $p$ -simplices, respectively.

More specifically,  $c_{i i'}^{(p-1)} = \|\alpha(\tau_i^{(p-1)}) - \alpha(\tau_{i'}^{(p-1)})\|_2$ , measures the Euclidean distance between the affine weight vectors of the  $i$ -th  $(p-1)$ -simplex of  $M_p$  and the  $i'$ -th  $(p-1)$ -simplex of  $M'_p$ , while  $c_{k k'}^{(p)} = \|\alpha(\sigma_k^{(p)}) - \alpha(\sigma_{k'}^{(p)})\|_2$ , measures the



**Fig. 1.** Pair of spheres with simplicial complexes constructed between the landmark points on the spheres along with assignments between cliques.

Euclidean distance between the affine weight vectors of the  $k$ -th  $p$ -simplex of  $M_p$  and the  $k'$ -th  $p$ -simplex of  $M'_p$ .

Similar to the *affinity matrix* construction in [38], we combine both the cost matrices in a single geodesic-cost matrix  $\mathcal{L}^{(p)} = (l_{ii',jj'}) \in \mathbb{R}^{nn' \times mm'}$  as

$$l_{ii',jj'}^{(p)} = \begin{cases} c_{ii'}^{(p-1)} & i = j, i' = j' \\ c_{kk'}^{(p)} & i \neq j, i' \neq j', a_{ik}a_{jk}a'_{i'k'}a'_{j'k'} = 1 \\ 0 & \text{otherwise} \end{cases}$$

The *diagonal* and *off-diagonal* entries of matrix  $\mathcal{L}^{(p)}$  capture the *Euclidean distances between the affine weight vectors of  $(p-1)$ -simplices* and the *Euclidean distances between the affine weight vectors of  $p$ -simplices*, respectively. Therefore, our QAP can now be formulated as

$$\begin{aligned} \operatorname{argmin}_{X_1, \dots, X_h} & \sum_{p=1}^h \operatorname{vec}(X_p)^T \mathcal{L}^{(p)} \operatorname{vec}(X_p) \\ \text{subject to} & \forall p \leq h, \mathbf{1}^T X_p = \mathbf{1}, X_p^T \mathbf{1} = \mathbf{1} \end{aligned} \quad (2)$$

where  $X_p$  is a permutation matrix and  $\operatorname{vec}(X_p)$  is its vector representation. Our solution to Equation 2 is concisely outlined in Algorithm 1. As we solve a QAP from highest to lowest dimension  $p$ -skeleton, we track the  $(p-1)$ -simplices whose matchings are *induced* by higher order simplex matches. On finding  $(p-1)$ -simplices that have the lowest cost and cannot be improved by solving a lower level QAP, we eliminate such simplices, causing the size of the matrix to shrink in subsequent iterations, leading to substantial speedups. Also, we use a *spectral relaxation* proposed by Lordeneu et. al. [20] to solve our QAP efficiently.

**Example** We illustrate with an example the bijective assignment produced by our algorithm between cliques / simplices of a pair of graph-induced *spherical* simplicial complexes, as shown in Figure 1. We consider two simplicial complexes

**Table 1.** Matchings of 3, 2-cliques of simplicial complexes  $\mathcal{K}$  and  $\mathcal{K}'$  shown in Figure 1.

$k$ -Clique	Matching between $\mathcal{K}$ & $\mathcal{K}'$	
3-Cliques	(1, 2, 3), (1, 3, 4), (3, 4, 5), (2, 18, 19), (5, 6, 7), (6, 8, 9), (12, 13, 14), (13, 14, 17), (14, 15, 16), (16, 19, 20).	(1, 2, 3), (1, 3, 4), (3, 4, 5), (2, 3, 19), (5, 6, 8), (6, 8, 9), (12, 13, 14), (13, 14, 17), (14, 15, 16), (16, 19, 20).
2-Cliques	(1, 2), (1, 3), (2, 3), (1, 4), (3, 4), (3, 5), (4, 5), (2, 19), (18, 19), (5, 6), (5, 7), (6, 8), (6, 9), (8, 9), (12, 13), (12, 14), (13, 14), (13, 17), (14, 17), (14, 15), (14, 16), (15, 16), (16, 19), (16, 20), (19, 20), (3, 8), (8, 18), (9, 10), (9, 13), (17, 18).	(1, 2), (1, 3), (2, 3), (1, 4), (3, 4), (3, 5), (4, 5), (2, 19), (3, 19), (5, 6), (5, 8), (6, 8), (6, 9), (8, 9), (12, 13), (12, 14), (13, 14), (13, 17), (14, 17), (14, 15), (14, 16), (15, 16), (16, 19), (16, 20), (19, 20), (3, 8), (8, 13), (9, 12), (9, 13), (17, 8).

$\mathcal{K}$  and  $\mathcal{K}'$  each embedded on  $\mathcal{S}^2$ , with 20 and 16 vertices, respectively. Matching of corresponding 3-cliques and 2-cliques are mentioned in the Table 1. Matching between vertices (1-cliques) is shown by marking them with the same label on both spheres.

**Time complexity analysis** The major cost incurred by our algorithm arises from matching cliques between two simplicial complexes. Therefore, we first derive an upper bound on the number of cliques that need to be matched as follows (proof in supplementary notes).

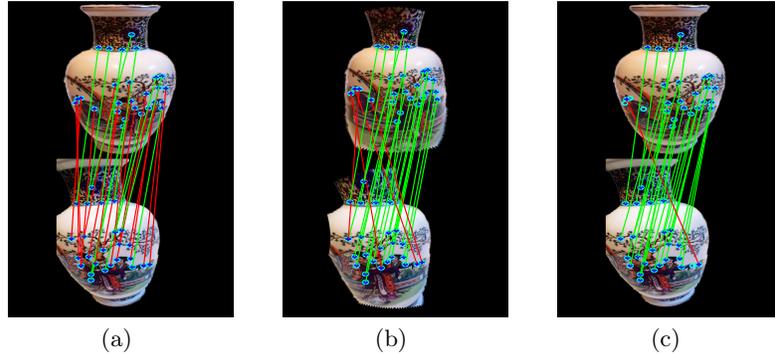
**Lemma 1.** *Let  $\mathcal{K}(G)$  represent the simplicial complex induced from graph  $G$  with  $n$  and  $m$  number of vertices and edges, respectively. Let  $h$  denote the maximum order of cliques in  $G$  and  $\delta$  be the maximum degree of a vertex in  $G$ . Then, the total number of  $k$ -cliques in  $\mathcal{K}(G)$  for  $k = (1, \dots, h)$ , are at most*

$$n + \frac{2m}{\delta(\delta+1)} \left[ \min \left\{ (\delta+1)^h + 1, \left( \frac{e(\delta+1)}{h} \right)^h \right\} - \delta - 2 \right]$$

Neglecting lower order terms, the number of cliques are of order  $O(n + m(\delta^{h-2} - \delta))$ . We know that the *spectral relaxation* proposed by Lordeneu et. al. [20] has a complexity of  $O(n^{3/2})$ , where  $n$  is the number of points to match on each side. Our higher order matching of cliques then has a time complexity of  $O(\{n + m(\delta^{h-2} - \delta)\}^{3/2})$ . In practice, for maximum order of cliques,  $h = 3$  (triangles) and  $h = 4$  (tetrahedrons), observe that the complexity drops to  $O(n^{3/2})$  and  $O(\{n + m\delta^2\}^{3/2})$ , respectively, which is very efficient.

## 5 Experiments

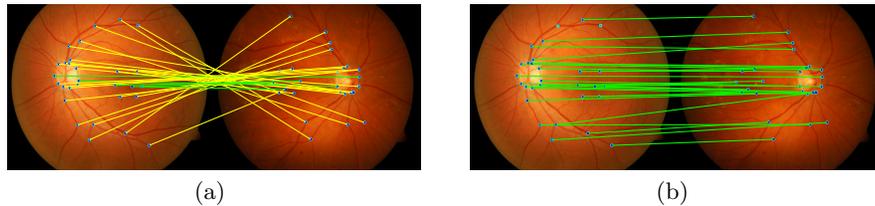
For our experiments, we considered *synthetic* and *real-world* datasets that cover both *spherical* and *planar* images. Spherical images can broadly be categorized as: *parabolic omnidirectional* ( $360^\circ$ ), *fish-eye*, and *panoramic* images. Note that our matching algorithm does not require any calibration parameters of cameras.



**Fig. 2.** Instances of matchings between (a) *Chinese vase* images for Tensor based method, (b) flat version of *Chinese vase* images for Tensor based method, and (c) *Chinese vase* images for our method. Green/red lines show correct/incorrect matches respectively. Isolated points show no matches.

To evaluate our matching algorithm, we compared against three main categories. (i) *Planar matching methods extended with geodesic metric on 3D manifolds*: Here, we *extended* the factorized graph matching (FGM) [38] algorithm by feeding it a  $k$ -NN graph based on geodesic distances between points to serve as our naive baseline method (called “FGM+geodesic”). The rest of the methods were feature-descriptor based. (ii) *Planar matching methods on 2D projected (unwrapped<sup>3</sup>) manifolds*. (iii) *Planar matching methods on 2D planar images*: Here, we proposed a *flat* version of our algorithm with Euclidean distance as the underlying metric (called “OurPlanar”) to work on flat 2D images.

Furthermore, we also perform experiments using RANSAC [16] for geometric verification and rectification. In our ablative studies, we analyze the robustness of our algorithm under affine transformations (rotation, reflection, scaling, and shear).



**Fig. 3.** Instances of matchings between (a) *Fundus* images for Tensor based method, (b) *Fundus* images for our method. Green/yellow lines show correct/incorrect matches respectively. Isolated points show no matches.

<sup>3</sup> *unwrapped*: planar projection of a spherical image with minimal distortion [9] [15].

**Baselines** We group the state-of-the-art methods as: (i) *Feature descriptor based matching for spherical and planar images*: BRISK [21], ORB [27], SPHORB [37]. (ii) *Graph based matching for planar images*: based on employing an *affinity matrix* (FGM) [38,39] and *eigenvalues* (EigenAlign) [13]. (iii) *Higher-order based matching for planar images*: using random clique complex (RCC) [30] and higher-order matching (Tensor) [12]. (iv) Finally, a naive baseline (FGM+geodesic) proposed by us that extends FGM by constructing a graph based on geodesic distances. Our code<sup>4</sup> is publicly available.

**Table 2.** Error (%) of pairwise matching between spherical images (omnidirectional, fish-eye and panorama) of five datasets for different methods.

Algorithms	Kamaishi	Chessboard	Desktop	Parking	Table
<i>OurWarped</i>	<b>0.79±0.0%</b>	<b>3.89±0.0%</b>	<b>0.32±0.0%</b>	<b>0.0±0.0%</b>	<b>0.74±0.0%</b>
FGM+geo	55.6±0.10%	79.2±1.21%	23.3±0.03%	37.5±0.0%	64.3±6.58%
SPHORB	90.0±0.0%	58.5±0.0%	91.1±0.0%	95.0±0.0%	78.5±0.0%
BRISK	85.6±0.0%	53.6±0.0%	78.9±0.0%	81.6±0.0%	69.2±0.0%
ORB	90.2±0.0%	53.8±0.0%	51.7±0.0%	71.1±0.0%	64.4±0.0%
Tensor	37.7±0.69%	60.5±0.41%	23.9±1.7%	23.7±7.5%	85.1±1.05%
FGM	53.3±0.21%	80.0±0.11%	31.9±0.12%	36.0±1.5%	65.5±0.01%

**Our method vs. planar matching methods with geodesic metric on 3D curved manifolds** Here, we match pairwise images directly on the warped images on curved manifolds (as shown in Figures 2, 3, and 4(b)). The comparison between standard higher-order graph matching (Tensor) [12] and our method on manifold is shown in Figures 2 and 3 using Chinese vases<sup>5</sup> and Fundus images [17], respectively. We observe from Figures 2(a) and 3(a) that the Tensor based method does not perform well on warped images. Although, the matching does improve when images are flattened to reduce the effect of curvature in Figure 2(b). Our method outperforms the baseline and has a maximum number of correct matches in Figures 2(c) and 3(b).

The error percentages of our warped image matching algorithm (OurWarped) are shown in the first row of Table 2. We observe that our method outperforms all other matching methods, including spherical feature descriptor based ones as well. Additional multimodal warped-planar matching experiments can be found in our supplementary notes.

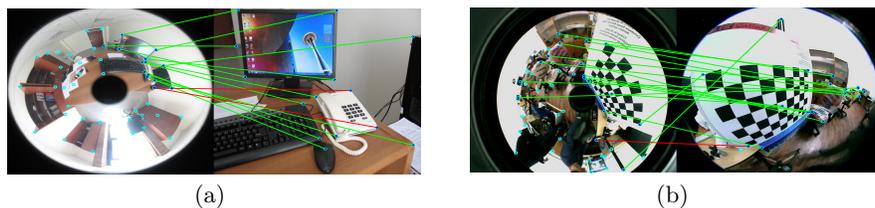
For matches between spherical and planar images, we find two variants which match between a spherical and a planar image (Figure 4(a)) and matching between different types of spherical images (Figure 4(b)). In Table 2, there is a slight increase in error percentages when matching across different types of spherical images, i.e., 3.89% for Chessboard, as compared to matching similar types, i.e., 0.32% for Desktop, due to differences in distortion levels. In spite of

<sup>4</sup> [Our Method](#) <sup>5</sup> from Google images

this, we find that our method significantly outperforms naive baseline and other matching methods on spherical images.

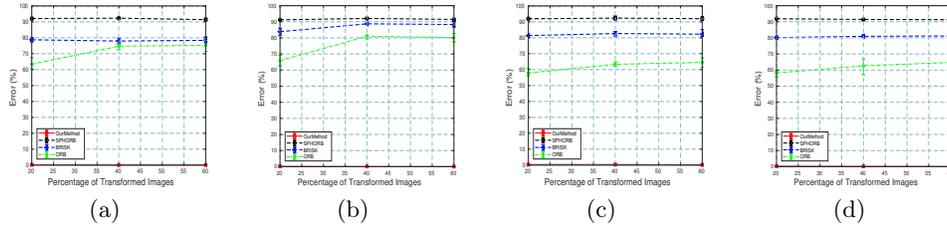
**Table 3.** Error (%) of pairwise matching between unwrapped equirectangular version of spherical (omnidirectional and fish-eye) images of four datasets for different methods including graph matching methods on flat surfaces.

Algorithms	Chessboard	Desktop	Parking	Table
<i>OurWarped</i>	<b>3.64 ± 0.0 %</b>	<b>1.06 ± 0.0 %</b>	<b>0.0 ± 0.0 %</b>	<b>0.57 ± 0.0 %</b>
RCC	28.6 ± 0.94 %	11.6 ± 0.74 %	13.2 ± 11.8 %	11.6 ± 0.57 %
EigenAlign	98.47 ± 0.0 %	95.24 ± 0.0 %	97.5 ± 0.0 %	97.9 ± 0.0 %
Tensor	68.9 ± 0.16 %	26.1 ± 0.58 %	19.0 ± 3.75 %	72.4 ± 0.67 %
FGM	84.0 ± 0.0 %	31.0 ± 0.0 %	38.0 ± 0.0 %	52.0 ± 0.0 %
SPHORB	58.6 ± 0.0 %	90.3 ± 0.0 %	97.5 ± 0.0 %	79.2 ± 0.0 %
BRISK	54.9 ± 0.0 %	84.9 ± 0.0 %	100.0 ± 0.0 %	74.2 ± 0.0 %
ORB	49.5 ± 0.0 %	78.2 ± 0.0 %	82.5 ± 0.0 %	70.3 ± 0.0 %



**Fig. 4.** Instances of matchings between (a) *Desktop* omnidirectional and planar images and (b) *Chessboard* omnidirectional and fish-eye images. Green/red lines show correct/incorrect matches, respectively. Isolated points show no matches.

***Our method vs. planar matching methods on 2D-projected curved manifolds*** Matching between spherical images can also be performed by applying planar graph matching methods on unwrapped equirectangular versions of spherical images. This makes the image flat and standard planar matching algorithms can then be employed. However, any kind of projection (on a flat surface in this case) introduces distortions in the resulting image. We flattened spherical images for four datasets mentioned in Table 3. We used two different methods to flatten omnidirectional and fish-eye images. The 360° image is unwrapped by dividing it into four parts (quadrants) and concatenated into a single flat image. On the other hand, fish-eye images do not cover the complete view of the scene and add distortion to the image due to curved mirrors and lenses of the cameras. We try to reduce the distortion by removing curves and flattening the image using calibration techniques outlined in [28]. Since any projection will lead to distortion,



**Fig. 5.** Error(%) in matching when varying the percentage (20% to 60%) of transformed images in the set of spherical images of *Desktop* (a)-(d). (a) 40° rotation, (b) reflection, (c) scaling and (d) shear.

we can compare the results from Table 2 with Table 3. Both the experimental outcomes are based on the same set of spherical images. Our matching algorithm significantly outperforms its competitors on both the spherical images and on curved manifolds.

**Our method vs. planar matching methods on 2D planar images** Our proposed method can also be reduced to a higher-order planar graph matching method. To show the importance of higher-order combinatorial matching not only with geodesic neighborhood, but also with euclidean neighborhood, we run our planar variant (OurPlanar) on popular 2D image datasets, competing with standard matching algorithms. We pick four well-known difficult matching datasets (Books, Building, Magazine, and Butterfly) that suffer from heavy occlusions and non-affine transformations [30]. Results for such an experiment are shown in Table 4. From the results, we observe that our method also serves as a powerful planar matching method and is still competitive using an euclidean neighborhood for our affine weight vectors. It significantly outperforms both the popular planar matching methods.

**Table 4.** Error (%) of pairwise matching between planar images of four datasets for different methods.

Algorithms	Magazine	Building	Books	Butterfly
<i>OurPlanar</i>	<b>0.0 ± 0.0 %</b>	<b>1.03 ± 0.01 %</b>	<b>19.72 ± 0.20 %</b>	<b>0.0 ± 0.0 %</b>
FGM	<b>0.0 ± 0.0 %</b>	74.87 ± 0.07 %	97.54 ± 0.01 %	16.12 ± 0.0 %
Tensor	<b>0.0 ± 0.0 %</b>	43.24 ± 2.98 %	32.35 ± 0.15 %	1.07 ± 0.17%

**RANSAC: Geometric Verification and Rectification** We also performed fundamental matrix based geometric verification using RANSAC algorithm [16] after descriptor based matching on two datasets for spherical and their planar versions for SPHORB and ORB in Table 5. We observed that the results are

improved (but still not better than our proposed method) in some cases but prune a lot of matches. Nearly 40 – 50% of matches are considered as outliers which makes it difficult to handle the noise. On the other hand, our method performs much better in any case while considering outliers.

**Table 5.** Error (%) of pairwise matching between spherical images of Desktop and Parking datasets and on their unwrapped versions for verification.

<b>Algorithms</b>	<b>Desktop</b>	<b>Desktop_flat</b>	<b>Parking</b>	<b>Parking_flat</b>
OurWarped	<b>0.32 ± 0.0 %</b>	<b>1.06 ± 0.0 %</b>	<b>0.0 ± 0.0 %</b>	<b>0.0 ± 0.0 %</b>
SPHORB+RANSAC	96.1 ± 0.0 %	93.9 ± 0.0 %	95.0 ± 0.0 %	100.0 ± 0.0 %
ORB+RANSAC	29.3 ± 0.0 %	70.6 ± 0.0 %	55.0 ± 0.0 %	100.0 ± 0.0 %

We performed rectification [5] on spherical images of Desktop dataset followed by BRISK descriptor for matching. The results improved from 78.9% (in Table 2) to 52.11% error. However, we observed that despite these improvements, our method still outperforms them. Also, in most of the cases, the rectification algorithm does not perform well and outputs noisy or distorted images. So, there is no guarantee to find the best solution.

***Ablative Studies (Effect of Affine Transformation)*** We remove completely at random 40 – 80% of landmark points on the Desktop dataset, and introduce affine transformations on these points. Figure 5 shows the results of affine transformation like *rotation*, *reflection*, *scaling*, and *shear*. We rotated images (clockwise) by 40° and performed matching for four algorithms. Then, we generated mirror images along the  $x$ -axis from the same dataset to introduce reflection. We also conducted transformation by scaling and shear of 360° images. We resized images in both the directions with scales 0.5 and 1.5 randomly. For shearing, we stretched images with 0.5 factor along  $y$ -axis. For all types of transformations, we observe that the results shown in Figure 5 clearly indicates that our method is robust to all kinds of affine transformations and easily outperforms other state-of-the-art methods.

## 6 Conclusion

We presented a bijective assignment between sets of landmark points embedded on a pair of images warped onto curved manifolds by the following steps. First, we built a *graph induced simplicial complex* on the warped images. Second, we proposed a constrained QAP that matches corresponding co-dimensional simplices between two simplicial complexes along with an efficient algorithm to solve the constrained QAP. Finally, we conducted extensive experiments, broadly grouped as *comparative matching* and *ablative studies*, in order to gain insight into the accuracy and robustness of our method. We are currently exploring the possibility of integrating such high-dimensional combinatorial structures into *Spherical CNNs* [6] to capture higher-order and latent structure.

## References

1. Ahmed, A.I., Aldhaheri, S.R., Rodriguez Kovacs, J., Narasimhulu, D., Putra, M., Minkoff, H., Haberman, S.: Sonographic measurement of cervical volume in pregnant women at high risk of preterm birth using a geometric formula for a frustum versus 3-dimensional automated virtual organ computer-aided analysis. *Journal of Ultrasound in Medicine* **36**(11), 2209–2217 (2017)
2. Alahi, A., Ortiz, R., Vandergheynst, P.: FREAK: Fast retina keypoint. In: CVPR. pp. 510–517. IEEE (2012)
3. Azevedo, R.G.d.A., Birkbeck, N., De Simone, F., Janatra, I., Adsumilli, B., Frossard, P.: Visual distortions in 360-degree videos. arXiv preprint arXiv:1901.01848 (2019)
4. Bay, H., Tuytelaars, T., Van Gool, L.: SURF: Speeded up robust features. In: ECCV. pp. 404–417. Springer (2006)
5. Bradski, G., Kaehler, A.: *Learning OpenCV: Computer vision with the OpenCV library*. O'Reilly Media Inc. (2008)
6. Cohen, T.S., Geiger, M., Köhler, J., Welling, M.: Spherical CNNs. ICLR (2018)
7. Colombo, C., Bimbo, A.D., Pernici, F.: Image mosaicing from uncalibrated views of a surface of revolution. In: Proc. BMVC. pp. 43.1–43.10 (2004)
8. Coors, B., Paul Condurache, A., Geiger, A.: Spherenet: Learning spherical representations for detection and classification in omnidirectional images. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 518–533 (2018)
9. Cruz-Mota, J., Bogdanova, I., Paquier, B., Bierlaire, M., Thiran, J.P.: Scale invariant feature transform on the sphere: Theory and applications. *International Journal of Computer Vision* **98**(2), 217–241 (2012)
10. Cyganek, B.: *An Introduction to 3D Computer Vision Techniques and Algorithms*. John Wiley & Sons, Inc. (2007)
11. Dey, T.K., Fan, F., Wang, Y.: Graph induced complex on point data. *Comput. Geom. Theory Appl.* **48**(8), 575–588 (2015)
12. Duchenne, O., Bach, F., Kweon, I.S., Ponce, J.: A tensor-based algorithm for high-order graph matching. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)* **33**(12), 2383–2395 (2011)
13. Feizi, S., Quon, G., Recamonde-Mendoza, M., Médard, M., Kellis, M., Jadbabaie, A.: Spectral alignment of networks. arXiv preprint arXiv:1602.04181 (2016)
14. Goesele, M., Snavely, N., Curless, B., Hoppe, H., Seitz, S.M.: Multi-view stereo for community photo collections. In: ICCV. pp. 1–8. IEEE (2007)
15. Guan, H., Smith, W.A.: BRISKS: Binary features for spherical images on a geodesic grid
16. Hartley, R., Zisserman, A.: *Multiple view geometry in computer vision*. Cambridge university press (2003)
17. Kalesnykiene, V., Kamarainen, J.k., Lensu, L., Sorri, I., Uusitalo, H., Kälviäinen, H., Pietilä, J.: Diaretdb0: evaluation database and methodology for diabetic retinopathy algorithms
18. Kaminsky, R.S., Snavely, N., Seitz, S.M., Szeliski, R.: Alignment of 3d point clouds to overhead images. In: 2009 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops. pp. 63–70. IEEE (2009)
19. Leordeanu, M., Zanfir, A., Sminchisescu, C.: Semi-supervised Learning and Optimization for Hypergraph Matching. In: IEEE International Conference on Computer Vision (ICCV) (November 2011)

20. Leordeanu, M., Hebert, M.: A spectral technique for correspondence problems using pairwise constraints. In: ICCV. vol. 2, pp. 1482–1489. IEEE (2005)
21. Leutenegger, S., Chli, M., Siegwart, R.Y.: BRISK: Binary robust invariant scalable keypoints. In: ICCV. pp. 2548–2555. IEEE (2011)
22. Li, Y., Gregori, G., Lam, B.L., Rosenfeld, P.J.: Automatic montage of sd-oct data sets. *Opt. Express* **19**(27), 26239–26248 (2011)
23. Lowe, D.G.: Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision* **60**(2), 91–110 (2004)
24. Munkres, J.R.: *Elements of algebraic topology*. Addison-Wesley (1984)
25. Olivier Pauly, Gozde Unal, G.S.S.C.T.F.: Semi-automatic matching of oct and ivus images for image fusion (2008), <https://doi.org/10.1117/12.773805>
26. Rosten, E., Drummond, T.: Machine learning for high-speed corner detection. In: ECCV. pp. 430–443. Springer (2006)
27. Rublee, E., Rabaud, V., Konolige, K., Bradski, G.: ORB: An efficient alternative to SIFT or SURF. In: ICCV. pp. 2564–2571. IEEE (2011)
28. Scaramuzza, D., Martinelli, A., Siegwart, R.: A toolbox for easily calibrating omnidirectional cameras. In: 2006 IEEE/RSJ International Conference on Intelligent Robots and Systems. pp. 5695–5701. IEEE (2006)
29. Sengupta, S., Singh, A., Leopold, H.A., Lakshminarayanan, V.: Ophthalmic diagnosis and deep learning – a survey (2018)
30. Sharma, C., Nathani, D., Kaul, M.: Solving partial assignment problems using random clique complexes. In: ICML. pp. 4593–4602 (2018)
31. Shen, D., Davatzikos, C.: HAMMER: Hierarchical attribute matching mechanism for elastic registration. *IEEE transactions on medical imaging* **21**(11), 1421–1439 (2002)
32. Starck, J., Hilton, A.: Spherical matching for temporal correspondence of non-rigid surfaces. In: Tenth IEEE International Conference on Computer Vision (ICCV’05) Volume 1. vol. 2, pp. 1387–1394. IEEE (2005)
33. Sung, B., Kim, S.H., Lee, S., Lim, J., Lee, J.K., Soh, K.S.: Nanofluid transport in a living soft microtube. *Journal of Physics D: Applied Physics* (2015)
34. Szeliski, R.: *Computer vision: algorithms and applications*. Springer Science & Business Media (2010)
35. Yang, J., Li, H., Jia, Y.: Go-icp: Solving 3d registration efficiently and globally optimally. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 1457–1464 (2013)
36. Zeng, A., Song, S., Nießner, M., Fisher, M., Xiao, J., Funkhouser, T.: 3dmatch: Learning local geometric descriptors from rgb-d reconstructions. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1802–1811 (2017)
37. Zhao, Q., Feng, W., Wan, L., Zhang, J.: SPHORB: A fast and robust binary feature on the sphere. *International Journal of Computer Vision* **113**(2), 143–159 (2015)
38. Zhou, F., De la Torre, F.: Factorized Graph Matching. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)* **38**(9), 1774–1789 (2016)
39. Zhou, F., De la Torre, F.: Deformable graph matching. In: CVPR. pp. 2922–2929 (2013)