

Feature matching using quasi-conformal maps*

Chun-xue WANG, Li-gang LIU^{†‡}

(School of Mathematical Sciences, University of Science and Technology of China, Hefei 230026, China)

[†]E-mail: lgliu@ustc.edu.cn

Received Nov. 26, 2015; Revision accepted Mar. 24, 2016; Crosschecked Apr. 22, 2017

Abstract: We present a fully automatic method for finding geometrically consistent correspondences while discarding outliers from the candidate point matches in two images. Given a set of candidate matches provided by scale-invariant feature transform (SIFT) descriptors, which may contain many outliers, our goal is to select a subset of these matches retaining much more geometric information constructed by a mapping searched in the space of all diffeomorphisms. This problem can be formulated as a constrained optimization involving both the Beltrami coefficient (BC) term and quasi-conformal map, and solved by an efficient iterative algorithm based on the variable splitting method. In each iteration, we solve two subproblems, namely a linear system and linearly constrained convex quadratic programming. Our algorithm is simple and robust to outliers. We show that our algorithm enables producing more correct correspondences experimentally compared with state-of-the-art approaches.

Key words: Feature correspondence; Quasi-conformal map; Splitting method

<http://dx.doi.org/10.1631/FITEE.1500411>

CLC number: TP391

1 Introduction

Finding geometrically consistent correspondences between pairs of images is one of the most fundamental operations in graphics and vision and it forms the basis of some problems such as feature tracking, image classification or retrieval, object detection, and shape matching. Given two images of the objects in different poses or even different objects of the same class, the task is to select a set of corresponding points, one from each image, such that the two points in a pair correspond to the same location. Identifying corresponding points is a challenging problem, as the shape of objects and their pose can significantly change across images.

Different matching strategies have been used among different image cues. Over the past decade, correspondence techniques have evolved significantly

in the graphics and vision literature. We refer the interested reader to Heider *et al.* (2011) and Tuytelaars and Mikolajczyk (2008) for a survey of the state-of-the-art methods in feature matching. These methods, such as the scale-invariant feature transform (SIFT) algorithm, create a collection of candidate correspondences by matching local signatures. However, as they consider only local intensity, many methods are globally inconsistent.

To filter the geometrically inconsistent correspondences, a low-dimensional parameterized model of deformations was proposed, such as in Chui and Rangarajan (2003), which performs feature-based nonrigid registration with the thin-plate spline (TPS). However, the resulting mapping may contain many incorrect pairs and only a few correct ones, as there is no guarantee that it is bijective. Focusing on the quality of the mapping, Lipman *et al.* (2014) selected a large subset of correspondences that are aligned by a global deformation of bounded distortion with an alternative approach, but the subset would filter more correct pairs while selecting more incorrect ones as the bounded distortion constraints

[‡] Corresponding author

* Project supported by the National Natural Science Foundation of China (Nos. 61672482 and 11626253) and the One Hundred Talent Project of the Chinese Academy of Sciences

 ORCID: Chun-xue WANG, <http://orcid.org/0000-0002-2118-3016>

© Zhejiang University and Springer-Verlag Berlin Heidelberg 2017

are set to be same everywhere. In this study, by balancing between these approaches, we focus mainly on the space of all diffeomorphisms by optimizing the deformation distance and keeping bijectivity.

Given several suitable landmarks, most existing algorithms can find registration in the space of all diffeomorphisms accurately and efficiently only by fixing the boundary vertex onto squares or circles (Zeng and Gu, 2011; Lui *et al.*, 2012; Lam and Lui, 2014). However, the registration problem becomes challenging when all points are set to be landmarks in the feature-matching problem. In this case, bijectivity is hard to keep in the obtained registration as the initial correspondences contain many outliers. To tackle this problem, we formulate an optimization problem of certain energy functionals over the space of all diffeomorphisms. Specifically, we minimize an energy functional involving both a Beltrami coefficient (BC) term and a deformation term, where the first term measures the distortion of the quasi-conformal map and the second measures the geometric distance between corresponding pairs. Different from Lipman *et al.* (2014), we find an orientation-preserving diffeomorphism $\tilde{f} : \{\mathbf{p}_i\}_{i=1}^N \rightarrow \{\mathbf{q}_i\}_{i=1}^N$ from a low-dimensional deformation space, where $\{\mathbf{p}_i\}_{i=1}^N$ and $\{\mathbf{q}_i\}_{i=1}^N$ are two input correspondences to be matched.

We propose an efficient splitting algorithm to perform this minimization in which we optimize variables alternatively. Using the proposed algorithm, diffeomorphisms (1-1 and onto) between two input feature data can be effectively obtained, even with the large number of incorrect matches. Numerical results show that our algorithm performs well in a number of experiments of synthetic data and real images over state-of-the-art approaches. The contributions of our work can be summarized as follows:

1. We are the first to introduce quasi-conformal maps into the feature-matching problem, which is formulated as a constrained optimization involving both the BCs and quasi-conformal map and solved by an efficient iterative algorithm based on variable splitting. Experiments show that our method can select more geometrically consistent correspondences compared with state-of-the-art approaches, even when the input contains many outliers.

2. Our method is not sensitive to the choice of the parameters and is robust to outliers. We demonstrate it on a series of synthetic data.

2 Related work

Finding meaningful mapping or matching between corresponding data that optimizes certain kinds of energy functionals has been extensively studied (Montagnat *et al.*, 2001; Nealen *et al.*, 2006; van Kaick *et al.*, 2011; Zhao *et al.*, 2012). In this section, we review several works on matching point features of nonrigid deformations. We also discuss some previous works closely related to registration via quasi-conformal maps.

1. Low-dimensional deformation spaces. With a set of candidate pairs of correspondences as input, which may include many outliers, several works aim to find a subset belonging to some low-dimensional deformation space \mathcal{D} . RANSAC (Fischler and Bolles, 1981) finds a large subset of pairs aligned by a global deformation in \mathcal{D} up to an ε -deviation denoted as \mathcal{D}_ε , which is formulated precisely as

$$\min_{f \in \mathcal{D}} \sum_{i=1}^N \|f(\mathbf{p}_i) - \mathbf{q}_i\|^0.$$

Lipman *et al.* (2014) performed a similar formulation to find the maximal subset of bounded distortion denoted as \mathcal{F}_K , and the problem is stated as follows:

$$\min_{f \in \mathcal{F}_K} \sum_{i=1}^N \|f(\mathbf{p}_i) - \mathbf{q}_i\|^0,$$

where \mathcal{F}_K is the collection of all deformations decided by a distortion value less than or equal to K , and both ε and K decide the deformation space to be low- or large-dimensional. Similarly, our method produces a low-dimensional deformation space represented by an optimal BC, denoted as \mathcal{D}_b .

2. Space-deformation based methods. Space deformation has been widely used in feature matching, and maximizes the quality of the matching and minimizes the complexity of the deformation simultaneously. Chui and Rangarajan (2003) minimized the deformation energy based on thin-plate splines, while Belongie *et al.* (2002) used shape context local descriptors to optimize a similar energy. Hinton *et al.* (1991) minimized deformation energy including a spline-based deformation cost and a generative model of appearance with an elastic matching algorithm. Jian *et al.* (2005) performed a nonrigid registration between these local frequency maps using the Riesz transform. However, these methods may be sensitive to initialization and outliers.

3. Quasi-conformal maps. We review mainly the work related to ours, including conformal maps, quasi-conformal maps, and Beltrami flows. Various conformal geometric methods have been presented for nonrigid surface matching and registration (Yezzi and Mennucci, 2005; Wang *et al.*, 2007; Zeng *et al.*, 2009). Taimouri and Hua (2014) introduced a new quasi-conformal metric which measures the curvature changes at each vertex of each pose during the deformation. Given boundary constraints, quasi-conformal maps are widely used to parameterize meshes onto disk domains and obtained by solving the Beltrami equation. Several methods have been proposed to deal with simple domains in the complex plane (Mastin and Thompson, 1984; Daripa, 1991; 1992; Lui and Ng, 2015). Ho and Lui (2016) proposed an algorithm to compute the quasi-conformal parameterization for a connected 2D domain or surface. A method called the Beltrami holomorphic flow was used to obtain the optimal BC associated with the registration (Lui *et al.*, 2012). Matching landmarks consistently with quasi-conformal mapping has also been proposed. Given several features on the surfaces, Zeng and Gu (2011) proposed a method to register 3D surfaces with large deformations using quasi-conformal curvature flow. However, bijectivity of the mapping is difficult to guarantee, especially when all feature pairs are set to be landmarks. Our approach can provide a diffeomorphism by setting each energy term a weight parameter properly.

3 Background

In this work, we apply quasi-conformal (QC) maps to find geometrically consistent correspondences between pairs of images. In this section, we review some basic theories on quasi-conformal geometry. We refer the interested reader to Gardiner and Lakic (2000) and Lehto *et al.* (1973) for details.

Given two surfaces M and N , a map $f: M \rightarrow N$ is conformal if it preserves the surface metric up to a multiplicative factor called the conformal factor. A generalization of conformal maps is the quasi-conformal maps, which are orientation-preserving homeomorphisms between Riemann surfaces with bounded conformality distortion, in the sense that their first-order approximations take small circles to small ellipses of bounded eccentricity (Gardiner and Lakic, 2000). Suppose $f: \mathbb{D} \rightarrow \mathbb{D}'$, where \mathbb{D} and \mathbb{D}'

are two domains in \mathbb{C} . Then f is quasi-conformal, provided that it satisfies the Beltrami equation

$$\frac{\partial f}{\partial \bar{z}} = \mu(z) \frac{\partial f}{\partial z} \quad (1)$$

for some complex-valued Lebesgue measurable $\mu: \mathbb{C} \rightarrow \mathbb{C}$, satisfying $\|\mu\|_\infty < 1$ (Bers, 1977). μ is called the Beltrami coefficient, which measures how far the map at each point deviates from a conformal map. Eq. (1) admits a geometrical interpretation. Equip \mathbb{D} with the metric tensor

$$ds^2 = \Omega(z)^2 |dz + \mu(z)d\bar{z}|^2,$$

where $\Omega(z) > 0$. Then, f satisfies Eq. (1) precisely when it is a conformal transformation from \mathbb{D} equipped with this metric to domain \mathbb{D}' equipped with the standard Euclidean metric, and f is called the μ -conformal. In particular, map f is conformal around the neighborhood of z when $\mu(z) = 0$.

Consider the effect of the pullback under f of the usual Euclidean metric. The resulting metric is then given by

$$\left| \frac{\partial f}{\partial \bar{z}} \right|^2 |dz + \mu(z)d\bar{z}|^2,$$

which is relevant to the background Euclidean metric $dzd\bar{z}$, and has eigenvalues

$$(1 + |\mu|)^2 \left| \frac{\partial f}{\partial \bar{z}} \right|^2,$$

and

$$(1 - |\mu|)^2 \left| \frac{\partial f}{\partial \bar{z}} \right|^2.$$

The eigenvalues represent the squared length of the major and minor axes of the ellipse obtained by pulling back along f the unit circle in the tangent plane (Fig. 1). Accordingly, the dilatation of f at a point z is defined by

$$K(z) = \frac{1 + |\mu(z)|}{1 - |\mu(z)|}.$$

The dilatation of f is given by

$$K(f) = \sup_{z \in \mathbb{D}} |K(z)| = \frac{1 + \|\mu\|_\infty}{1 - \|\mu\|_\infty}.$$

A simple calculation leads to the following property:

Theorem 1 If $f: \mathbb{C} \rightarrow \mathbb{C}$ satisfies $\|\mu(f)\|_\infty < 1$, then f is a diffeomorphism (Gardiner and Lakic, 2000).

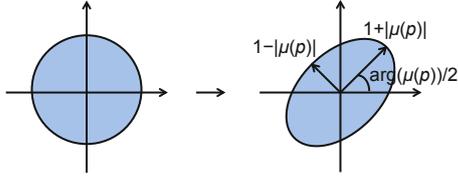


Fig. 1 Illustration of how the Beltrami coefficient controls conformality distortion

Theorem 1 plays the fundamental role of obtaining a diffeomorphism. That is, given a BC $\mu : \mathbb{C} \rightarrow \mathbb{C}$ with $\|\mu\|_\infty < 1$, we can find a one-to-one quasi-conformal mapping from \mathbb{C} onto itself satisfying the Beltrami equation (Eq. (1)).

4 Problem formulation

Given a set of candidate pairs of correspondences in the complex plane $(\mathbf{p}_i, \mathbf{q}_i) \in \mathbb{C} \times \mathbb{C}$ ($i = 1, 2, \dots, N$), our aim is to extract a geometrically consistent subset $\{(\mathbf{p}_l, \mathbf{q}_l)\}$, $l = 1, 2, \dots, n \leq N$ matched by $\tilde{f} \in \mathcal{D}_b$ (\mathcal{D}_b is the space of all diffeomorphisms). We can set up our problem as

$$\tilde{f} = \arg \min_{f \in \mathcal{D}_b} \sum_{i=1}^N \|f(\mathbf{p}_i) - \mathbf{q}_i\|^0,$$

where the ℓ_0 -norm measures the Euclidean distance between $f(\mathbf{p}_i)$ and \mathbf{q}_i ; i.e., $\|f(\mathbf{p}_i) - \mathbf{q}_i\|^0 = 1$ if $f(\mathbf{p}_i) \neq \mathbf{q}_i$, and $\|f(\mathbf{p}_i) - \mathbf{q}_i\|^0 = 0$ otherwise. In this way, we select matches with small conformal distortion. Unfortunately, it is impossible to construct space \mathcal{D}_b by adding several constraints, such as in Lipman *et al.* (2014). Going back to Eq. (1), we solve Beltrami's equation by involving an L^p -minimization of the Beltrami energy:

$$E_B(f) = \left\| \frac{\partial f}{\partial \bar{z}} - \mu \frac{\partial f}{\partial z} \right\|^p. \quad (2)$$

Note that, if we set p large enough, the optimization problem gives a good approximation. In particular, when $p = 2$, the result of the least-square problem (2) is called the least-square quasi-conformal map associated to μ (Gu and Yau, 2008; Weber *et al.*, 2012). In this study, we choose $p = 2$ and solve it with an alternating minimization algorithm.

Going back to Theorem 1, to obtain a diffeomorphic registration associated with an optimal BC, which is guaranteed to be bijective, our optimization can generally be written as finding $\tilde{f} : S_1 \rightarrow S_2$ that

satisfies

$$\begin{aligned} \tilde{f} &= \arg \min_f \mathbf{E}(f) \\ &= \int |\nabla \mu_f|^2 dS + \alpha \int |\mu_f|^2 dS \\ &\quad + \gamma \int |f_{\bar{z}} - \mu_f f_z|^2 dS + \beta \sum_{i=1}^N \|f(\mathbf{p}_i) - \mathbf{q}_i\|^0 \end{aligned} \quad (3)$$

$$\text{s.t. } \|\mu_f\|_\infty < 1,$$

where μ_f is the BC associated with $f \in \mathcal{D}_b$, and we write \int_{S_1} as \int for short.

The first term of \mathbf{E} is a regularization term ensuring the smoothness of f . The second term of \mathbf{E} is minimized to control the conformality distortion of f . The third term of \mathbf{E} aims to minimize the discretization of the least-square Beltrami energy, and the last term is designed to align as many of \mathbf{p}_i 's as we can with \mathbf{q}_i 's.

The optimization problem (3) poses two main challenges:

1. The problem involves the optimization of both quasi-conformal maps f and μ_f , which depend highly on each other, leading to a highly nonlinear optimization.
2. The $\mathcal{L}_{2,0}$ functional in problem (3) is not smooth and is nonconvex.

We tackle the problem by combining the splitting method and iterative reweighted least-square (IRLS) method. More specifically, following Lipman *et al.* (2014), we replace the $\mathcal{L}_{2,0}$ term by a smooth functional to approximate its minima. Since μ_f and f are highly related and the energy is highly nonlinear, we introduce a new variable ν and reformulate problem (3) to be an equality-constrained problem. Thus, the problem can be further formulated as looking for an optimal BC $\tilde{\nu} : S_1 \rightarrow \mathbb{C}$, which is the BC of some diffeomorphism $\tilde{f} : S_1 \rightarrow S_2$, minimizing the following energy functional $\mathbf{E}(\nu, f)$:

$$\begin{aligned} (\tilde{\nu}, \tilde{f}) &= \arg \min_{f, \nu} \mathbf{E}(\nu, f) \\ &= \int |\nabla \nu|^2 dS + \alpha \int |\nu|^2 dS + \gamma \int |f_{\bar{z}} - \nu f_z|^2 dS \\ &\quad + \beta \sum_{i=1}^N \|f(\mathbf{p}_i) - \mathbf{q}_i\|^0 \\ \text{s.t. } \|\nu\|_\infty &< 1, \nu = \mu(f). \end{aligned} \quad (4)$$

Problems (3) and (4) are equivalent to each other. Based on problem (4), we apply the splitting method

and the IRLS method to solve the constrained optimization problem, by replacing the $\mathcal{L}_{2,0}$ functional term and equation constraints, as follows:

$$\begin{aligned}
 (\tilde{\nu}, \tilde{f}) &= \arg \min_{f, \nu} \mathbf{E}_\delta^{\text{split}}(\nu, f) \\
 &= \int |\nabla \nu|^2 dS + \alpha \int |\nu|^2 dS \\
 &\quad + \sigma \int |\nu - \mu(f)|^2 dS + \gamma \int |f_{\bar{z}} - \nu f_z|^2 dS \quad (5) \\
 &\quad + \beta \sum_{i=1}^N (\|f(\mathbf{p}_i) - \mathbf{q}_i\|^2 + \delta)^{\frac{p}{2}} \\
 \text{s.t. } &\|\nu\|_\infty < 1,
 \end{aligned}$$

where p is a small constant, the sequence $\delta \rightarrow 0$ is a parameter updated during the iterations, and σ is the weight of the penalty term required to increase to infinity theoretically. We iteratively minimize $\mathbf{E}_\delta^{\text{split}}(\nu, f)$ subject to the constraints. More precisely, given an initial match $f^{(0)}$, set $\nu^{(0)} = 0$, $\delta = \delta^{(0)}$. Suppose we have obtained $\nu^{(k)}$ and $\delta^{(k)}$ at the k th iteration. Fixing $\nu^{(k)}$, we obtain $f^{(k)}$ by minimizing $\mathbf{E}_{\delta^{(k)}}^{\text{split}}(\nu^{(k)}, f)$ over f . Once $f^{(k)}$ is updated, similarly, by fixing $f^{(k)}$, we obtain $\nu^{(k+1)}$ by minimizing $\mathbf{E}_{\delta^{(k)}}^{\text{split}}(\nu, f^{(k)})$ over ν . If $|\mathbf{E}_{\delta^{(k+1)}}^{\text{split}}(\nu^{(k+1)}, f^{(k+1)}) - \mathbf{E}_{\delta^{(k)}}^{\text{split}}(\nu^{(k)}, f^{(k)})| < \varepsilon$, update $\delta^{(k)}$.

5 Numerical implementation

5.1 Discretization

We use a Delaunay triangulation of the planar point set $\{\mathbf{p}_i\}$ denoted as $T = (V, F)$, where $V = \{\mathbf{p}_i\}_{i=1}^N$ is the set of vertices and $F = \{\mathbf{f}_j\}_{j=1}^T$ is the set of oriented faces. We choose f from the space of continuous piecewise linear (CPL) mappings defined by the values at vertices. The complex derivatives $f_{\bar{z}}$ and f_z are naturally defined on each triangle. Given $\mathbf{e}_i, \mathbf{e}_j, \mathbf{e}_k$ as edge vectors of T with opposite vertices (i, j, k) , the gradient of f is $(f_i \mathbf{t}_i + f_j \mathbf{t}_j + f_k \mathbf{t}_k) / (2A_T)$, where A_T is area of the triangle, $\mathbf{t}_i = \mathbf{e}_i^\perp$, and f_i is the value at vertex \mathbf{p}_i . Then per-triangle derivatives are given by

$$\begin{aligned}
 f_z &= \frac{1}{4A_T} (f_i \bar{t}_i + f_j \bar{t}_j + f_k \bar{t}_k), \\
 f_{\bar{z}} &= \frac{1}{4A_T} (f_i t_i + f_j t_j + f_k t_k),
 \end{aligned}$$

where $t_i = t_i^x + i t_i^y$ is the complex form of $\mathbf{t}_i = (t_i^x, t_i^y)$.

The BC μ is a piecewise constant function defined on each triangle of T . Different from Lam and Lui (2014) and Lui and Ng (2015), who computed μ on each vertex and converted the value to each face by averaging the values on each vertex, we restrict the value of μ on triangle τ as μ_τ and compute it by $\mu_\tau = f_{\bar{z}}|_\tau / f_z|_\tau$ (defined above) for each triangle, so $\boldsymbol{\mu} = (\mu_0, \mu_1, \dots, \mu_{T-1}) \in \mathbb{R}^T$. For any $\boldsymbol{\mu}^1, \boldsymbol{\mu}^2, \boldsymbol{\mu} \in \mathbb{R}^T$, we define the inner product and norm as

$$\begin{aligned}
 (\boldsymbol{\mu}^1, \boldsymbol{\mu}^2)_{\mathbb{R}^T} &= \sum_{\tau} \mu_\tau^1 \mu_\tau^2 A_\tau, \\
 \|\boldsymbol{\mu}\|_{\mathbb{R}^T} &= \sqrt{(\boldsymbol{\mu}, \boldsymbol{\mu})_{\mathbb{R}^T}}.
 \end{aligned}$$

For any $\boldsymbol{\mu} \in \mathbb{R}^T$, we define the jump of $\boldsymbol{\mu}$ over an edge e as

$$[\boldsymbol{\mu}]_e = \begin{cases} \sum_{e \prec \tau} \mu_\tau \text{sgn}(e, \tau), & e \not\subseteq \partial T, \\ 0, & e \subseteq \partial T, \end{cases}$$

where $e \prec \tau$ denotes that e is an edge of triangle τ , and $\text{sgn}(e, \tau)$ defines the relative orientation of edge e to triangle τ . More specifically, we set all the triangles with an anticlockwise orientation, and all edges are with fixed orientations chosen randomly. For an edge e , if the orientation of e is consistent with the orientation of τ , then $\text{sgn}(e, \tau) = 1$; otherwise, $\text{sgn}(e, \tau) = -1$. In the piecewise constant function space, the gradient operator can be defined as

$$\nabla : \mu \rightarrow \nabla \mu, \quad \nabla \mu_e = [\boldsymbol{\mu}]_e.$$

Thus, the term $\int |\nabla \mu|^2 dS$ can be discretized as $\sum_e l_e [\boldsymbol{\mu}]_e^2$, where l_e denotes the length of edge e . Since $\|\mu_\tau\|_\infty < 1$ is hard to optimize, we have simplified it into $\|\mu_\tau^x\|_\infty < \sqrt{2}/2$ and $\|\mu_\tau^y\|_\infty < \sqrt{2}/2$ for all $\tau \in F$.

To summarize, we discretize energy problem (5) with complex variables f_i to present the map and per-face complex variables ν_τ for BC:

$$\begin{aligned}
 (\tilde{\nu}, \tilde{f}) &= \arg \min_{f, \nu} \mathbf{E}^{\text{split}}(\nu, f) \\
 &= \frac{1}{n_e} \sum_e l_e [\boldsymbol{\nu}]_e^2 + \frac{\alpha}{T} \sum_\tau |\nu_\tau|^2 + \frac{\sigma}{T} \sum_\tau A_\tau |\nu_\tau - \mu_\tau(f)|^2 \\
 &\quad + \frac{\gamma}{T} \sum_\tau A_\tau |f_{\bar{z}}|_\tau - \nu_\tau f_z|_\tau|^2 \\
 &\quad + \frac{\beta}{N} \sum_{i=1}^N (\|f(\mathbf{p}_i) - \mathbf{q}_i\|^2 + \delta)^{\frac{p}{2}} \\
 \text{s.t. } &\|\nu_\tau^j\|_\infty < \frac{\sqrt{2}}{2}, \quad \tau \in F, \quad j = 0, 1,
 \end{aligned} \tag{6}$$

where n_e denotes the number of edges, and A_τ the area of triangle τ .

5.2 Optimization

Energy problem (6) is highly nonlinear. Standard methods, e.g., Newton methods or ADMM methods (Boyd *et al.*, 2011; Li *et al.*, 2015), do not perform well for optimizing this energy. Besides, it is quadratic with respect to both vector variables $\mathbf{f} = [f(\mathbf{p}_1), f(\mathbf{p}_2), \dots, f(\mathbf{p}_N)]$ and $\boldsymbol{\nu} = [\nu_1, \nu_2, \dots, \nu_T]$. This suggests the use of an alternating-descent algorithm to optimize \mathbf{f} and $\boldsymbol{\nu}$ in an alternating fashion. More specifically, if \mathbf{f} is fixed, we can obtain $\boldsymbol{\nu}$ by solving a linearly constrained convex quadratic problem. Similarly, we can easily obtain \mathbf{f} by solving a linear system when $\boldsymbol{\nu}$ is fixed. While there is no guarantee of the convergence to a global minimum, this algorithm is relatively stable with consistent practical behavior (Fig. 7).

1. **f**-subproblem. We discuss the minimization of $\mathbf{E}_{p,\delta}^{\text{split}}(\boldsymbol{\nu}^{(k)}, \mathbf{f})$ over \mathbf{f} by fixing $\boldsymbol{\nu}^{(k)}$. Note that for each $p > 0$ and $\delta > 0$, the standard family of functionals

$$\mathbf{E}_{p,\delta}(\mathbf{f}) = \sum_{i=1}^N (\|f(\mathbf{p}_i) - \mathbf{q}_i\|^2 + \delta)^{\frac{p}{2}}$$

is smooth, and we can take a sequence $\delta \rightarrow 0$ while p is a small constant. Let us denote $\mathbf{E}_\delta^{(k)}$ as the energy at the k th iteration:

$$\mathbf{E}_\delta^{(k)} = \sum_{i=1}^N \omega_i^{(k-1)} \|f(\mathbf{p}_i) - \mathbf{q}_i\|^2. \quad (7)$$

Since $\mu(\mathbf{f})$ characterizes the distortion of the map (scale, stretch, etc.), we set a small weight to the region of large distortion, that is,

$$\omega_i^{(k-1)} = \exp\left(-\frac{\|\mu(\mathbf{f}^{(k-1)})_i^v\|^2}{2m}\right) \cdot d_i, \quad (8)$$

where

$$d_i = \left(\|f^{(k-1)}(\mathbf{p}_i) - \mathbf{q}_i\|^2 + \delta^{(k)}\right)^{\frac{p}{2}-1},$$

$\mu(\cdot)_i^v$ denotes the value of $\mu(\cdot)$ on the i th vertex computed by averaging the values of its neighboring faces, and we set $m = 0.005$ in our implementation.

Then the energy is given by

$$\begin{aligned} \mathbf{f}^{(k)} = \arg \min_{\mathbf{f}} & \left\{ \frac{\gamma}{T} \sum_{\tau} A_{\tau} \left| f_{\bar{z}}|_{\tau} - \nu_{\tau}^{(k)} f_z|_{\tau} \right|^2 \right. \\ & \left. + \frac{\beta}{N} \sum_{i=1}^N \omega_i^{(k-1)} \|f(\mathbf{p}_i) - \mathbf{q}_i\|^2 \right\}, \end{aligned} \quad (9)$$

which is a quadratic problem in \mathbf{f} whose solution can be easily obtained by solving a linear system.

2. **ν**-subproblem. Once $\mathbf{f}^{(k)}$ is obtained, we minimize $\mathbf{E}_{p,\delta}^{\text{split}}(\boldsymbol{\nu}, \mathbf{f}^{(k)})$ over $\boldsymbol{\nu}$ while fixing $\mathbf{f}^{(k)}$. That is, we obtain $\boldsymbol{\nu}^{(k+1)}$ by solving the following problem:

$$\begin{aligned} \min_{\boldsymbol{\nu}} & \mathbf{E}^{\text{split}}(\boldsymbol{\nu}, \mathbf{f}^{(k)}) \\ & = \frac{1}{n_e} \sum_e l_e [\boldsymbol{\nu}]_e^2 + \frac{\sigma}{T} \sum_{\tau} A_{\tau} \left| \nu_{\tau} - \mu_{\tau}(\mathbf{f}^{(k)}) \right|^2 \\ & \quad + \frac{\alpha}{T} \sum_{\tau} A_{\tau} |\nu_{\tau}|^2 + \frac{\gamma}{T} \sum_{\tau} \left| f_{\bar{z}}^{(k)}|_{\tau} - \nu_{\tau} f_z^{(k)}|_{\tau} \right|^2 \\ \text{s.t.} & \|\nu_{\tau}^j\|_{\infty} < \frac{\sqrt{2}}{2}, \quad \tau \in F, \quad j = 0, 1, \end{aligned} \quad (10)$$

where $\mu(\mathbf{f})$ is the BC associated with \mathbf{f} .

Theoretically, the conventional penalty method requires that σ should increase to infinity. In our experiments, we found that setting σ to be a large constant can also give satisfactory results. Obviously, this is a linearly constrained convex quadratic programming problem, optimized over $\boldsymbol{\nu}$ with Matlab's quadratic programming routine. The two-step optimization is summarized in Algorithm 1.

6 Experimental results

We tested our algorithm on both synthetic data and real images, and compared it with several existing algorithms, including BD (Lipman *et al.*, 2014), Tensor (Duchenne *et al.*, 2011), and RANSAC (Fischler and Bolles, 1981) algorithms, both visually and quantitatively. For comparison, we tested the Tensor method with a cost function to penalize changes in the angles between triplets of feature points and compute three RANSAC models, affine (denoted as RANSAC-AFF), epipolar (denoted as RANSAC-EPI), and projective (denoted as RANSAC-PRO) (using Matlab's `estimateGeometricTransform` and `estimateFundamentalMatrix` function, respectively), with suitable parameters demonstrated in Lipman *et al.* (2014). In our experiments we chose constants $p = 0.001$ and $\sigma = 20$. We decreased δ by half in

Algorithm 1 Feature matching using quasi-conformal maps

Require: candidate correspondences $(\mathbf{p}_i, \mathbf{q}_i) \in \mathbb{R}^2 \times \mathbb{R}^2$, $i = 1, 2, \dots, N$

Ensure: subset of pairs $(\mathbf{p}_{i_l}, \mathbf{q}_{i_l})$, $l = 1, 2, \dots, n \leq N$, and the bijective map \mathbf{f}

- 1: **Initialization:** Triangulation $T = \{V, F\}$ from $\{\mathbf{p}_i\}_{i=1}^N$
 $\mathbf{f}^{(0)}(\mathbf{p}_i) = \mathbf{p}_i$, $\boldsymbol{\nu}^{(0)} = \mathbf{0}$
 $\delta^{(0)} \approx \text{diam}\{\mathbf{p}_i\}$, ω_i^0 from Eq. (8)
// $\text{diam}\{\cdot\}$ gives the minimum radius of a circle that contains all the points \mathbf{p}_i
 $k = 0$, $\delta_{\min} = 1 \times 10^{-8}$, $k_{\max} = 50$
- 2: **while** $\delta^{(k)} > \delta_{\min}$ and $k < k_{\max}$ **do**
- 3: **while** $E_{\delta^{(k)}}^{\text{split}}(\boldsymbol{\nu}^{(k)}, \mathbf{f}^{(k)}) - E_{\delta^{(k-1)}}^{\text{split}}(\boldsymbol{\nu}^{(k-1)}, \mathbf{f}^{(k-1)}) > \varepsilon$ **do**
- 4: $k = k + 1$
- 5: $\delta^{(k)} = \delta^{(k-1)}$
- 6: Set $\omega_i^{(k-1)}$ according to Eq. (8)
- 7: Optimize \mathbf{f} subproblem via problem (9) with $\boldsymbol{\nu}^{(k)}$ fixed
- 8: Optimize $\boldsymbol{\nu}$ subproblem via problem (10) with $\mathbf{f}^{(k)}$ fixed
- 9: **end while**
- 10: $\delta^{(k)} = \delta^{(k)} / 2$
- 11: **end while**
- 12: Return all pairs $(\mathbf{p}_i, \mathbf{q}_i)$ for $(\mathbf{p}_i, \mathbf{q}_i) < \varepsilon$

each iteration, and set the lower bound as 1×10^{-8} . Experiments showed that our algorithm is not sensitive to these choices. To measure whether $\mathbf{f}(\mathbf{p}_i)$ reaches the location closely to the target \mathbf{q}_i , we use the distance threshold of five pixels.

6.1 Synthetic data

To compare our algorithm with other state-of-the-art methods, we first tested it on synthetic data. We sampled points of a square randomly and uniformly on the 2D plane, and produced another set of points by a random radial basis function (RBF) mapping of the first one. We produced six RBFs. For each RBF, we produced $n = 64$ inlier pairs of points sampled uniformly from an 8×8 grid $\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_n$, chose 10 control points randomly denoted as $\{\mathbf{c}_k\}_{k=1}^{10}$, and then obtained the corresponding RBF as

$$\mathbf{f}(\mathbf{p}) = \sum_{i=1}^{10} \phi(d(\mathbf{p}, \mathbf{p}_i)) \mathbf{c}_i + \mathbf{P}_1(\mathbf{p}),$$

where $\phi(\cdot)$ is an RBF, and we used $\phi(r) = 1 - 30r^2 - 10r^3 + 45r^4 - 6r^5 - 60r^3 \cdot \log r$. \mathbf{P}_1 is a vector valued

linear polynomial. Next, following Lipman *et al.* (2014), $N - n$ outlier pairs $(\mathbf{p}_i, \mathbf{q}_i)$ were added by first creating \mathbf{p}_i uniformly in the square, $i = n + 1, n + 2, \dots, N$, and then choosing \mathbf{q}_i according to the distribution of outliers estimated from real images.

To evaluate the performance of our algorithm, we carried out 100 trials randomly with an outlier fraction in the range $[0, 0.95]$ for each RBF and computed the F -measure for each trial:

$$F = 2 \cdot \frac{\text{precision} \cdot \text{recall}}{\text{precision} + \text{recall}},$$

where precision is the fraction of retrieved instances that are relevant, and recall is the fraction of relevant instances that are retrieved (Sasaki, 2007). We plotted the F -measure of the six RBFs as a function of the fraction of outliers over all trials in Fig. 2, which shows that our algorithm gives a higher F -measure even for a large fraction of outliers.

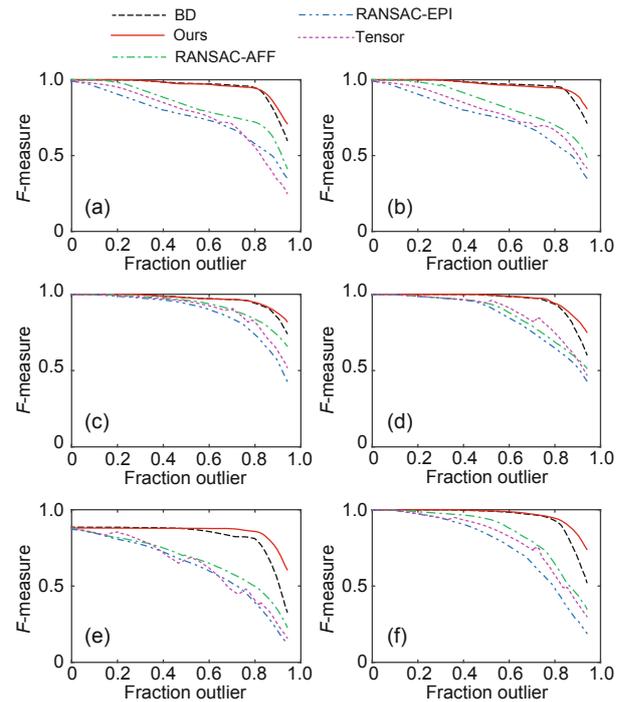


Fig. 2 Performance curves on six RBFs using our method and other methods: (a) RBF 1; (b) RBF 2; (c) RBF 3; (d) RBF 4; (e) RBF 5; (f) RBF 6

6.2 Real images

Results on real images are shown in Figs. 3–5. We tested our algorithm on different kinds of image pairs including images with large deformations, images from different viewpoints, and images

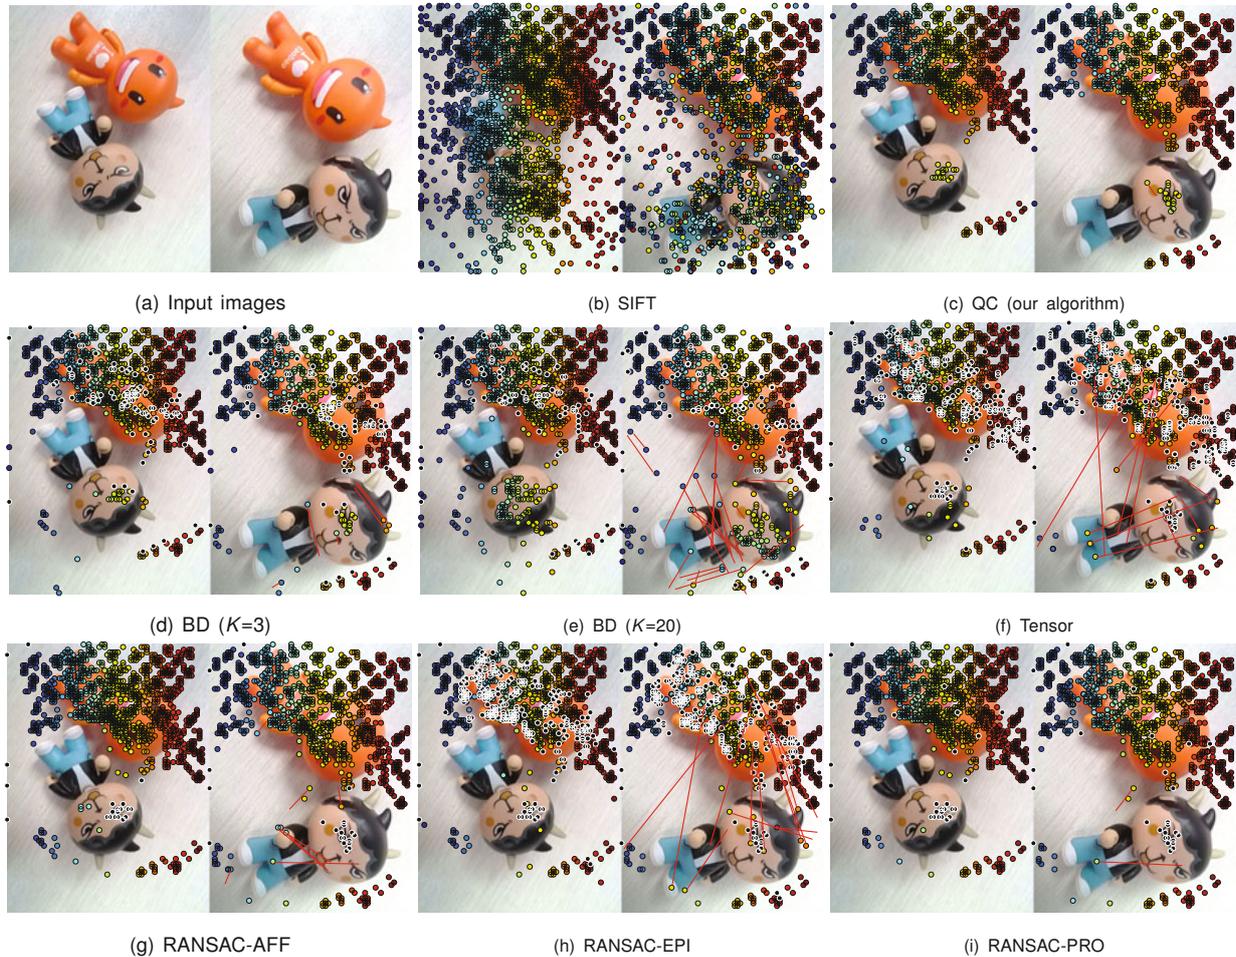


Fig. 3 Experiments on real images with large deformations. There are nine image pairs: original images (a), SIFT correspondences as input (b), and several methods for filtering the SIFT correspondences (c)–(i). Disks of the same color denote correspondences in an image pair, and the color varies linearly according to the horizontal position of points in the left image. On the right image of each pair, red lines show the errors between selected correspondences and ground-truth correspondences. The black disks are used to visualize the pairs selected by our method but missed by other methods (References to color refer to the online version of this figure)

of different animals of the same species. We used SIFT descriptors in the VLFeat software package to obtain the initial correspondences in all experiments (Vedaldi and Fulkerson, 2010). We compared RANSAC, Tensor, and BD algorithms with distortion bound parameter $K = 3$ and $K = 20$ characterizing different deformation spaces.

In Figs. 3–5, disks of the same size and color indicate matching pairs to visualize matching results clearly. In the left image of each image pair, we set the color varying according to the horizontal position of the points. For all real images, we used ground-truth correspondences marked by a human observer,

and showed the correspondence error by connecting a red line between each correspondence chosen and the respective ground-truth correspondence in each right image. Our goal is to select as many SIFT correspondences compatible with the ground-truth or with small positional deviations (indicated by short red lines), while filtering the ones with large positional deviations (indicated by long red lines). In addition, we used black disks to visualize the pairs selected by our method but missed by other methods. The results show that our method selects more consistent correspondences than the others.



Fig. 4 Experiments on real images from different viewpoints: red lines show the errors between selected correspondences and ground-truth correspondences and the black disks in the results are used to visualize the pairs selected by our method but missed by other methods (see caption in Fig. 3 for details). The images were taken with permission from Lazebnik *et al.* (2004; 2005) (References to color refer to the online version of this figure)



Fig. 5 Experiments on real images of different animals of the same species: red lines show the errors between selected correspondences and ground-truth correspondences and the black disks in the results are used to visualize the pairs selected by our method but missed by other methods (see caption in Fig. 3 for details). The images were taken with permission from Lazebnik *et al.* (2004; 2005) (References to color refer to the online version of this figure)

6.2.1 Images with large deformations

For images with large deformations, the SIFT descriptors are less discriminative and contain more incorrect correspondences, and it is difficult to filter correct correspondences. Fig. 3 shows a scene where two objects are close to each other in the left image while away from each other in the right one. In this case, the distances are preserved between feature points within each object while changing significantly across the two objects, and our aim is to

capture the large deformation while discarding the outliers. Fig. 3 shows that previous methods cannot capture correct pairs near the eyes since the input SIFT matches contain many incorrect ones. In contrast, our method selects more good correspondences even though the SIFT matches are of poor quality.

6.2.2 Images from different viewpoints

We tested our method on the images from different viewpoints (Fig. 4). It is a challenging test case for images from different viewpoints since some

feature points are captured in the first image, while they are in shade in the second one. In addition, the location of one object with respect to another or background may change significantly. Examples in Fig. 4 demonstrate that previous works either miss more good pairs (indicated by black disks) and introduce more false ones (indicated by long red lines) or fail to recognize symmetry and repetitions (RANSAC-AFF). Our method can capture repeated structures or breaks correctly and choose more pairs fitting the input with fewer outliers.

6.2.3 Images of different animals of the same species

Fig. 5 shows the performance of our algorithm over animal images from the natural world. We tested the images of different animals of the same species. Since the SIFT descriptor is unstable, many outliers are included, but our algorithm can still perform well. The first example in Fig. 5a shows that RANSAC-PRO and BD ($K = 3$) lose more correct correspondences and other algorithms always select more incorrect ones, while we can balance them well especially on the body of the bird. For the other example, both RANSAC-EPI and RANSAC-PRO lose more correct correspondences on the right wing, while RANSAC-AFF cannot filter the outliers efficiently. However, we can locate the left and right wings correctly with more consistent matches.

7 Discussion

7.1 Quality

To quantify the correspondences retrieved, we present the F -measure in Table 1 on the examples in Figs. 3–5 of the tested algorithms. For each example, since there might be some discrepancy between

a human observer and the geometrically consistent SIFT correspondences, we define ground-truth as an acceptable distance threshold of five pixels from the human observer. From Table 1, we can find that our algorithm achieves larger F -measures than the other algorithms.

7.2 Bijectivity

Following Section 3, we investigated the bijectivity of deformation f by computing the norm of the BC on each face. Fig. 6 shows the histograms of the BC norms of examples from Figs. 3–5. All the BC norms are less than 1.

7.3 Running time

We have implemented the alternating-descent algorithm using MATLAB on a laptop with Pentium IV, 2.16 GHz CPU, and 2 GB RAM. We solved f using a sparse linear system and ν using quadratic programming in Matlab. Table 2 gives the model statistic and running time of the examples from Figs. 3–5. Our method takes more time than RANSAC and BD algorithms, as it solves a quadratic program and a linear system in each iteration.

7.4 Convergence

We solve energy problem (6) by the alternating-descent algorithm discussed in Section 5. Although it is difficult to prove the convergence of the splitting algorithm, the energy decreases consistently through iterations, similar to the coordinate descent algorithms (Wright, 2015). We further examine the convergence of our algorithm via numerical experiments and plot the energy curves of examples in Figs. 3–5 (Fig. 7). The results show that the energy always decreases during the iterations.

Table 1 Comparison of F -measure of real images

Method	F -measure				
	Fig. 3	Figs. 4a–4i	Figs. 4j–4r	Figs. 5a–5i	Figs. 5j–5r
BD ($K = 3$)	0.919	0.943	0.796	0.896	0.847
BD ($K = 5$)	0.899	0.922	0.773	0.854	0.806
BD ($K = 20$)	0.835	0.808	0.564	0.714	0.728
RANSAC-AFF	0.910	0.938	0.773	0.862	0.786
RANSAC-EPI	0.851	0.848	0.782	0.705	0.561
RANSAC-PRO	0.924	0.890	0.733	0.902	0.776
Tensor	0.859	0.812	0.664	0.744	0.813
Ours	0.945	0.972	0.858	0.962	0.959

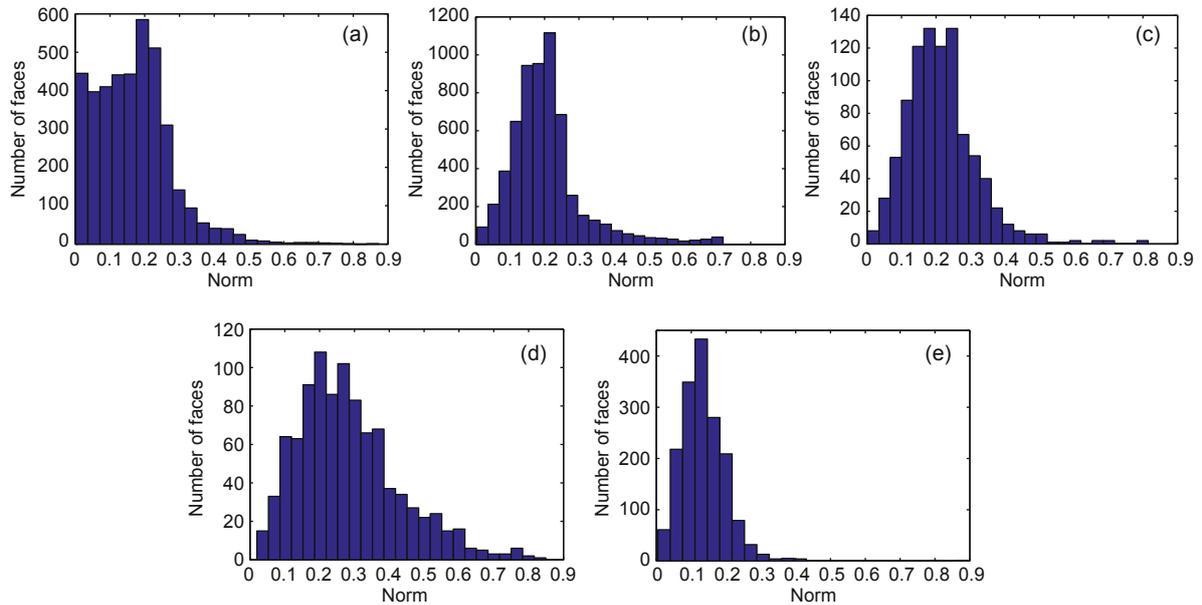


Fig. 6 Histograms of the BC norm of examples in Fig. 3 (a), Figs. 4a–4i (b), Figs. 4j–4r (c), Figs. 5a–5i (d), and Figs. 5j–5r (e)

Table 2 Comparison of running time

Method	Running time (s)				
	Fig. 3 ($ v =1905, f =3980$)	Figs. 4a–4i ($ v =2925, f =6064$)	Figs. 4j–4r ($ v =495, f =908$)	Figs. 5a–5i ($ v =533, f =980$)	Figs. 5j–5r ($ v =902, f =1690$)
BD	215.491	441.196	52.367	60.543	105.553
RANSAC-AFF	0.028	0.151	0.411	0.257	0.092
RANSAC-EPI	10.436	10.543	7.145	7.361	7.537
RANSAC-PRO	0.023	0.790	0.533	0.522	0.572
Tensor	792.466	1827.547	132.693	169.544	321.593
Ours	242.911	410.278	54.736	74.501	122.583

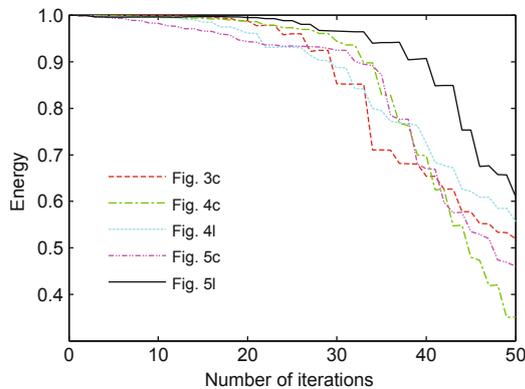


Fig. 7 Energy curves of examples in Figs. 3–5. These curves show that the energy decreases during the iterations and yields a stable result

7.5 Choice of parameters

We have tested our algorithm with different parameter settings on RBF 1. Specifically, we have applied our algorithm with different values of σ in the interval (10, 30) on data RBF 1 and ε in the interval (2, 8) to see the average performance with different outlier fractions of data RBF 1. Fig. 8 shows that our algorithm is not sensitive to all of these choices.

8 Conclusions

In this work, we considered the problem of finding a geometrically consistent set of correspondences between two input images. Given a set of candidate matches provided by SIFT descriptors, which may include many outliers, we selected a subset of these correspondences that can be aligned perfectly from

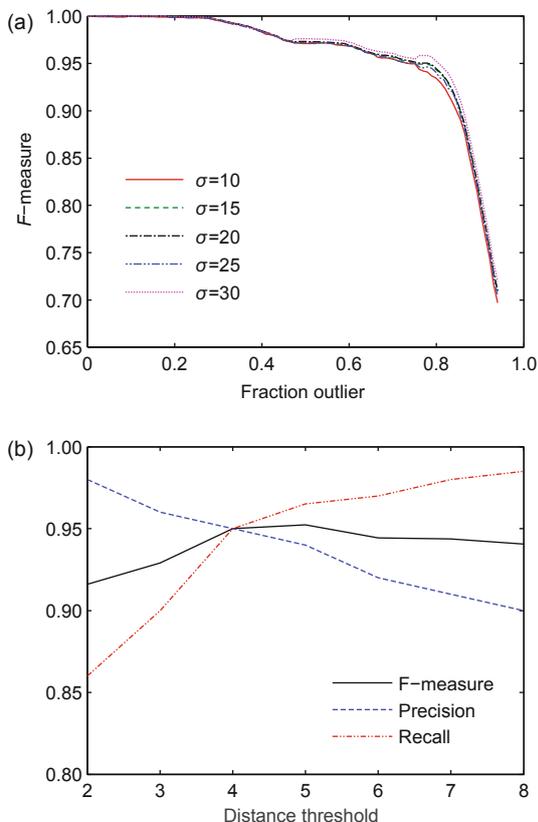


Fig. 8 The performance of our method with different choices of parameter σ (a) and the effect of varying the distance threshold ε (b). The performance is stable for a wide range of parameter values

the space of all diffeomorphisms. We formulated this problem as a constrained optimization problem involving both the BC term and quasi-conformal map term, and solved it using a splitting method. In each iteration of our algorithm, we solved the f subproblem by solving a linear system and ν subproblem by linearly constrained convex quadratic programming. We further examined the convergence of our algorithm via numerical experiments. Experiments showed that our algorithm is not sensitive to the parameter choice and is robust to outliers, and produces excellent results on synthetic data and real images in comparison to the state-of-the-art approaches.

We dealt with the one-to-one correspondence of geometrically consistent features in this work, but there are still many issues to be solved. Given two sets of features, which may contain different numbers of features, we hope that we can generalize this algorithm to find the overall consistent correspondence mapping (one-to-one or one-to-many). In addition,

since SIFT may not be stable when the deformations and lighting changes are large, how to build a descriptor capturing more geometric information would be an interesting topic for future work.

References

- Belongie, S., Malik, J., Puzicha, J., 2002. Shape matching and object recognition using shape contexts. *IEEE Trans. Patt. Anal. Mach. Intell.*, **24**(4):509-522. <http://dx.doi.org/10.1109/34.993558>
- Bers, L., 1977. Quasiconformal mappings, with applications to differential equations, function theory and topology. *Bull. Am. Math. Soc.*, **83**(6):1083-1100. <http://dx.doi.org/10.1090/S0002-9904-1977-14390-5>
- Boyd, S., Parikh, N., Chu, E., et al., 2011. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Found. Trends Mach. Learn.*, **3**(1):1-122. <http://dx.doi.org/10.1561/22000000016>
- Chui, H., Rangarajan, A., 2003. A new point matching algorithm for non-rigid registration. *Comput. Vis. Image Understand.*, **89**(2-3):114-141. [http://dx.doi.org/10.1016/S1077-3142\(03\)00009-2](http://dx.doi.org/10.1016/S1077-3142(03)00009-2)
- Daripa, P., 1991. On a numerical method for quasi-conformal grid generation. *J. Comput. Phys.*, **96**(1):229-236. [http://dx.doi.org/10.1016/0021-9991\(91\)90274-O](http://dx.doi.org/10.1016/0021-9991(91)90274-O)
- Daripa, P., 1992. A fast algorithm to solve nonhomogeneous Cauchy-Reimann equations in the complex plane. *SIAM J. Sci. Stat. Comput.*, **13**(6):1418-1432. <http://dx.doi.org/10.1137/0913080>
- Duchenne, O., Bach, F., Kweon, I.S., et al., 2011. A tensor-based algorithm for high-order graph matching. *IEEE Trans. Patt. Anal. Mach. Intell.*, **33**(12):2383-2395. <http://dx.doi.org/10.1109/TPAMI.2011.110>
- Fischler, M.A., Bolles, R.C., 1981. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM*, **24**(6):381-395. <http://dx.doi.org/10.1145/358669.358692>
- Gardiner, F.P., Lakic, N., 2000. Quasiconformal Teichmüller Theory. American Mathematical Society, Providence, USA. <http://dx.doi.org/10.1090/surv/076>
- Gu, X.D., Yau, S.T., 2008. Computational Conformal Geometry. International Press, Somerville, MA, USA.
- Heider, P., Pierre-Pierre, A., Li, R., et al., 2011. Local shape descriptors, a survey and evaluation. Eurographics Workshop on 3D Object Retrieval, p.1-8. <http://dx.doi.org/10.2312/3DOR/3DOR11/049-056>
- Hinton, G.E., Williams, C.K.L., Revow, M.D., 1991. Adaptive elastic models for hand-printed character recognition. 4th Int. Conf. on Neural Information Processing Systems, p.512-519.
- Ho, K.T., Lui, L.M., 2016. QCMC: quasi-conformal parameterizations for multiply-connected domains. *Adv. Comput. Math.*, **42**(2):279-312. <http://dx.doi.org/10.1007/s10444-015-9424-1>
- Jian, B., Vemuri, B.C., Marroquin, J.L., 2005. Robust nonrigid multimodal image registration using local frequency maps. Biennial Int. Conf. on Information Processing in Medical Imaging, p.504-515. http://dx.doi.org/10.1007/11505730_42

- Lam, K.C., Lui, L.M., 2014. Landmark and intensity-based registration with large deformations via quasi-conformal maps. *SIAM J. Imag. Sci.*, **7**(4):2364-2392. <http://dx.doi.org/10.1137/130943406>
- Lazebnik, S., Schmid, C., Ponce, J., 2004. Semi-local affine parts for object recognition. British Machine Vision Conf., p.779-788. <http://dx.doi.org/10.5244/C.18.98>
- Lazebnik, S., Schmid, C., Ponce, J., 2005. A maximum entropy framework for part-based texture and object recognition. ICCV, p.832-838. <http://dx.doi.org/10.1109/ICCV.2005.10>
- Lehto, O., Virtanen, K.I., Lucas, K.W., 1973. Quasiconformal Mappings in the Plane. Springer New York.
- Li, Y., Xie, X., Yang, Z., 2015. Alternating direction method of multipliers for solving dictionary learning. *Commun. Math. Stat.*, **3**:37-55. <http://dx.doi.org/10.1007/s40304-015-0050-5>
- Lipman, Y., Yagev, S., Poranne, R., et al., 2014. Feature matching with bounded distortion. *ACM Trans. Graph.*, **33**(3):26. <http://dx.doi.org/10.1145/2602142>
- Lui, L.M., Ng, T.C., 2015. A splitting method for diffeomorphism optimization problem using Beltrami coefficients. *J. Sci. Comput.*, **63**(2):573-611. <http://dx.doi.org/10.1007/s10915-014-9903-4>
- Lui, L.M., Wong, T.W., Zeng, W., et al., 2012. Optimization of surface registrations using Beltrami holomorphic flow. *J. Sci. Comput.*, **50**(3):557-585. <http://dx.doi.org/10.1007/s10915-011-9506-2>
- Mastin, C.W., Thompson, J.F., 1984. Quasiconformal mappings and grid generation. *SIAM J. Sci. Stat. Comput.*, **5**(2):305-310. <http://dx.doi.org/10.1137/0905022>
- Montagnat, J., Delingette, H., Ayache, N., 2001. A review of deformable surfaces: topology, geometry and deformation. *Image Vis. Comput.*, **19**(14):1023-1040. [http://dx.doi.org/10.1016/S0262-8856\(01\)00064-6](http://dx.doi.org/10.1016/S0262-8856(01)00064-6)
- Nealen, A., Müller, M., Keiser, R., et al., 2006. Physically based deformable models in computer graphics. *Comput. Graph. For.*, **25**(4):809-836. <http://dx.doi.org/10.1111/j.1467-8659.2006.01000.x>
- Sasaki, Y., 2007. The Truth of the F -measure. School of Computer Science, University of Manchester.
- Taimouri, V., Hua, J., 2014. Deformation similarity measurement in quasi-conformal shape space. *Graph. Models*, **76**(2):57-69. <http://dx.doi.org/10.1016/j.gmod.2013.12.001>
- Tuytelaars, T., Mikolajczyk, K., 2008. Local invariant feature detectors: a survey. *Found. Trends Comput. Graph. Vis.*, **3**(3):177-280. <http://dx.doi.org/10.1561/06000000017>
- van Kaick, O., Zhang, H., Hamarneh, G., et al., 2011. A survey on shape correspondence. *Comput. Graph. For.*, **30**(6):1681-1707. <http://dx.doi.org/10.1111/j.1467-8659.2011.01884.x>
- Vedaldi, A., Fulkerson, B., 2010. VLFeat: an open and portable library of computer vision algorithms. Proc. 18th ACM Int. Conf. on Multimedia, p.1469-1472. <http://dx.doi.org/10.1145/1873951.1874249>
- Wang, S., Wang, Y., Jin, M., et al., 2007. Conformal geometry and its applications on 3D shape matching, recognition, and stitching. *IEEE Trans. Patt. Anal. Mach. Intell.*, **29**(7):1209-1220. <http://dx.doi.org/10.1109/TPAMI.2007.1050>
- Weber, O., Myles, A., Zorin, D., 2012. Computing extremal quasiconformal maps. *Comput. Graph. For.*, **31**(5):1679-1689. <http://dx.doi.org/10.1111/j.1467-8659.2012.03173.x>
- Wright, S.J., 2015. Coordinate descent algorithms. *Math. Program.*, **151**(1):3-34. <http://dx.doi.org/10.1007/s10107-015-0892-3>
- Yezzi, A., Mennucci, A., 2005. Conformal metrics and true "gradient flows" for curves. ICCV, p.913-919. <http://dx.doi.org/10.1109/ICCV.2005.60>
- Zeng, W., Gu, X.D., 2011. Registration for 3D surfaces with large deformations using quasi-conformal curvature flow. CVPR, p.2457-2464. <http://dx.doi.org/10.1109/CVPR.2011.5995410>
- Zeng, W., Hua, J., Gu, X., 2009. Symmetric conformal mapping for surface matching and registration. *Int. J. CAD/CAM*, **9**(1):103-109.
- Zhao, Z., Feng, X., Teng, S., et al., 2012. Multiscale point correspondence using feature distribution and frequency domain alignment. *Math. Probl. Eng.*, **2012**:382369. <http://dx.doi.org/10.1155/2012/382369>