



HAL
open science

Higher-order Graph Principles towards Non-rigid Surface Registration

Yun Zeng, Chaohui Wang, Dimitris Samaras, Xianfeng Gu, Nikos Paragios

► **To cite this version:**

Yun Zeng, Chaohui Wang, Dimitris Samaras, Xianfeng Gu, Nikos Paragios. Higher-order Graph Principles towards Non-rigid Surface Registration. [Research Report] RR-8607, Inria Saclay Ile de France; INRIA. 2014, pp.31. hal-01086052v1

HAL Id: hal-01086052

<https://inria.hal.science/hal-01086052v1>

Submitted on 21 Nov 2014 (v1), last revised 30 Sep 2015 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Higher-order Graph Principles towards Non-rigid Surface Registration

Yun Zeng, Chaohui Wang, Xianfeng Gu, Dimitris Samaras, Nikos Paragios

**RESEARCH
REPORT**

N° 8607

September 2014

Project-Teams GALEN



Higher-order Graph Principles towards Non-rigid Surface Registration

Yun Zeng^{*}, Chaohui Wang[†], Xianfeng Gu^{*}, Dimitris Samaras^{*},
Nikos Paragios^{‡ §}

Project-Teams GALEN

Research Report n° 8607 — September 2014 — 28 pages

Abstract: This report casts surface registration as a problem of finding a set of discrete correspondences through the minimization of an energy function endowed with geometric/appearance similarities and higher-order deformation priors. Two higher-order graph-based formulations are proposed under different deformation assumptions. The first formulation encodes the isometric deformation based on conformal geometry in a higher-order graph matching problem, which is solved through dual-decomposition and is able to handle partial matching. Despite the isometry assumption, this approach is able to robustly match two sparse feature point sets on two surfaces even undergoing heavily anisometric deformation. Nevertheless, its performance degrades significantly when extended to performing anisometric registration for a set of densely sampled points. This issue is rigorously addressed subsequently through a deformation model that is able to handle arbitrary diffeomorphisms between two surfaces. Such a deformation model is then introduced into a higher-order MRF formulation for dense surface registration, which we solve with a new parallel and memory efficient algorithm. To deal with the prohibitive search space, we design an efficient way to select a number of matching candidates for each point of the source surface based on the matching results of a sparse set of points. Our method is validated through a series of experiments demonstrating its accuracy and efficiency, notably in challenging cases of large and/or anisometric deformations, or meshes that are partially occluded.

Key-words: 3D Surface Registration, Higher-order Graph Matching, Conformal Geometry, Markov Random Fields

^{*} Department of Computer Science, Stony Brook University, Stony Brook, NY, 11794, USA

[†] Laboratoire d'Informatique Gaspard Monge, UMR CNRS 8049, Université Paris-Est, 77454 Marne-la-Vallée Cedex 2, France

[‡] Center for Visual Computing, Ecole Centrale Paris

[§] Equipe GALEN, INRIA Saclay - Île de France, Orsay, France

**RESEARCH CENTRE
SACLAY – ÎLE-DE-FRANCE**

Parc Orsay Université
4 rue Jacques Monod
91893 Orsay Cedex

Principes de Graphes d'Ordre élevé pour le Recalage Non-Rigide de Surfaces 3D

Résumé : Ce rapport formalise le problème du recalage de surfaces 3D comme la recherche d'un ensemble de correspondances discrètes, par la minimisation d'une fonction d'énergie basée sur des similitudes géométriques et d'apparence, et des à priori d'ordre élevé sur la déformation. Deux formulations à base de graphes d'ordre élevé sont proposées sous différentes hypothèses de déformation. La première formulation encode la déformation isométrique, à partir de géométrie conforme, dans un problème d'appariement de graphes d'ordre élevé, qui est résolu par décomposition duale et est capable de gérer les cas de correspondance partielle. Malgré l'hypothèse d'isométrie, cette approche est capable de mettre en correspondance de manière robuste deux ensembles clairsemés de points sur deux surfaces, y compris lorsque celles-ci subissent une déformation fortement anisométrique. Cependant, sa performance se dégrade de manière significative lorsqu'elle est étendue au recalage anisométrique d'un ensemble de points à forte densité. Ce problème est rigoureusement traité par la suite à travers un modèle de déformation capable de gérer des difféomorphismes arbitraires entre deux surfaces. Ce modèle de déformation est introduit dans une formulation MRF d'ordre élevé pour le recalage dense de surfaces, que nous résolvons avec un nouvel algorithme parallèle et efficace en termes de mémoire. Pour traiter l'espace de recherche prohibitif, nous concevons une méthode efficace pour sélectionner un ensemble de correspondants potentiels pour chaque point appartenant à la surface source. Cette méthode est basée sur les résultats d'appariement d'un ensemble clairsemé de points. Notre méthode est validée au moyen d'une série d'expériences qui démontrent sa précision et son efficacité, notamment dans les cas difficiles où des déformations importantes et/ou anisométriques sont présentes, ou lorsque les maillages sont partiellement cachés.

Mots-clés : Recalage de Surfaces 3D, Appariement de Graphes d'ordre élevé, Géométrie Conforme, Champs de Markov Aléatoires d'ordre élevé

1 Introduction

The problem of finding the correspondences between two or more surfaces is the prerequisite for many applications in computer vision and medical imaging such as shape recognition, deformation transfer, and object recognition [19]. Furthermore, the proliferation of 3D content (*e.g.*, [1, 57, 74, 82]) has further enhanced the need of developing a robust surface registration method, in particular for 3D shapes undergoing large, non-rigid deformations and/or are acquired with a high level of noise.

Surface registration approaches can be classified as either extrinsic (where the coordinates for each point on the surface is known, *e.g.*, [8]), or intrinsic (where the surface is represented in a parametrization space, *e.g.*, [15]). In order to handle surfaces undergoing large deformations, directly searching in their extrinsic space is often intractable due to the large degrees of freedom in the non-rigid deformations of the shapes. In contrast, intrinsic space based methods can significantly reduce the complexity of the problem by representing the shape in a space that is invariant to certain types of deformations.

In the most general case, intrinsic space based methods assume that two surfaces undergo only *isometric* (*i.e.*, distance or metric preserving) deformations. Such an assumption can be a rough approximation for the deformations of most real-world objects. State-of-the-art intrinsic space based methods include geodesic/exponential maps ([13, 14]), conformal maps ([73, 77, 81]) and diffusion maps ([51, 58]). Among these approaches, only conformal maps provide a closed-form solution to the dense surface matching problem. The solution is characterized by the Möbius transformation in the uniformization domain, which requires the determination of up to three correspondences between two surfaces. Conformal maps can also be directly generalized to anisometric deformations as shown in [2], where a quasi-conformal mapping was able to handle arbitrary diffeomorphisms between two surfaces.

Although conformal maps are well-suited for establishing dense correspondences between surfaces undergoing large deformations, relying on a single mapping is often sensitive to noisy input, or boundary conditions. It is therefore necessary to consider the correspondences obtained from multiple mappings and define robust metrics to handle uncertainties in the input. To this purpose, graph-based approaches have been very successful in handling noisy input and occlusions [38, 70].

In this paper, we cast the surface registration problem into the search of the correspondences for a discrete set of points on the surface. We show that the deformation constraints between two surfaces can be represented by up to third-order interactions among these point sets. Thus, finding the optimal registration boils down into the determination of the optimal matching for a higher-order graph by minimizing an energy which constrains the quality of any possible matching. To efficiently integrate the powers of intrinsic space-based methods and graph-based approaches, we propose two graph-based formulations to solve the sparse and dense surface registration problems respectively, allowing us to achieve both efficiency and accurate dense surface registration results.

In the first stage, we formulate the surface registration problem as a higher-order graph matching problem that integrates both extrinsic matching costs (*e.g.*, texture, curvature and normal consistencies) and intrinsic deformation metric (deviation from isometry, inspired by [47]), which is solved using the dual-decomposition technique [42]. Despite its success in registering sparsely sampled points, the deformation metric [47] is only robust to points on the surface that are far apart. When the points to be registered get dense, any possible triplet matching would result in a small distance measure, leading the metric to be no longer discriminating.

Towards a dense surface registration, for the sake of model accuracy, we propose a novel local surface deformation model to characterize arbitrary deformations between two surfaces. Our new deformation model is derived using the ideas in Riemannian geometry [23] in which a shape is equipped with a metric tensor at each point in the parametrization domain. For any given point, we consider a special canonical parametrization of the shape whose metric tensor at that point is Euclidean, leading to a generic deformation descriptor, canonical distortion coefficients (CDCs), which can be efficiently computed in the discrete setting. With the above deformation model, we formulate the dense surface registration

into a higher-order MRF optimization problem. Furthermore, we design an efficient matching candidate selection scheme for any point on the surface based on the fact that any three correspondences determine a unique closed-form solution to establish the mapping between two surfaces, together with a parallel optimization algorithm, to efficiently solve such an MRF problem.

In summary, the main contributions of this paper include:

- A higher-order graph approach for the surface registration that deals with anisometric deformation between two surfaces.
- A novel surface deformation model for anisometric deformation and its use in a proposed higher-order MRF formulation for dense surface matching.
- An efficient implementation of the higher-order MRF optimization that requires minimal memory and achieves significant speedup by parallel acceleration.

Note that the idea of using multiple Möbius transformations induced by any three correspondences to achieve dense registration first appeared in the conference version of this paper [79], which is extended in later works such as [35].

This paper is organized as follows: In Sec. 1.1, we review related work on surface matching, graph matching and MRFs. An overview of our surface registration system is given in Sec. 1.2. Following this, we present in detail our sparse surface registration algorithm based on higher-order graph matching as well as its optimization in Sec. 2, and the dense surface matching algorithm based on higher-order MRFs also together with its optimization in Sec. 3. The implementation details and experimental results are given in Sec. 4 and we conclude the whole work in Sec. 5.

1.1 Related work

Modeling surface deformation is a very challenging task due to the large degrees of freedom that exist in real-world deformations. The “rigidity” assumption has been widely used, as a trade-off between the quality in deformation approximation and the simplicity in computation. Existing deformation models either characterize the *rigidity* in the extrinsic space or in the intrinsic space.

The choice of extrinsic representations of shapes in the presence of global rigid deformation (*i.e.*, rotation and translation) was handled in the literature exhaustively through the iterative closest points (ICP) [8] method with various extensions (*e.g.*, [16, 29]). However, global rigidity does not take into account bendable shapes (*e.g.*, garments or rubber bands). Hence the notion of *local rigidity* has been proposed endowing the assumption that the deformation between two local corresponding neighborhoods is rigid (*e.g.*, [30]). It should be noted though that searching for the correspondences between two surfaces with large deformations directly in the original extrinsic space may suffer from high computational complexity.

In a number of applications/computer vision tasks such as registration of subtle facial expressions, there are localized, high degrees of freedom deformations. Surface representations through distance functions and expression of the surface registration problem as a volume registration one were studied initially to provide a dense solution to the correspondence problem [31, 49, 52]. Such methods could handle a reasonable but still limited amount of deformations due to the regularization constraints imposed on the deformation field. To tackle large scale deformations, several approaches have been developed to obtain dense point correspondences by representing the surfaces to a canonical domain which preserves the geodesic distances or angles [13, 14, 73, 77, 81]. Such representations usually require an initial set of feature correspondences or boundary conditions. Given noisy 3D scan data with varying scales, boundaries and resolutions, the performance of the above methods is compromised from the fact that it is difficult to find reliable feature point correspondences and consistent boundary conditions. Furthermore, since most surface deformations are not perfectly isometric, solely considering intrinsic information introduces approximation errors to the matching result. Due to the above-mentioned issues, [5, 63, 64] proposed to

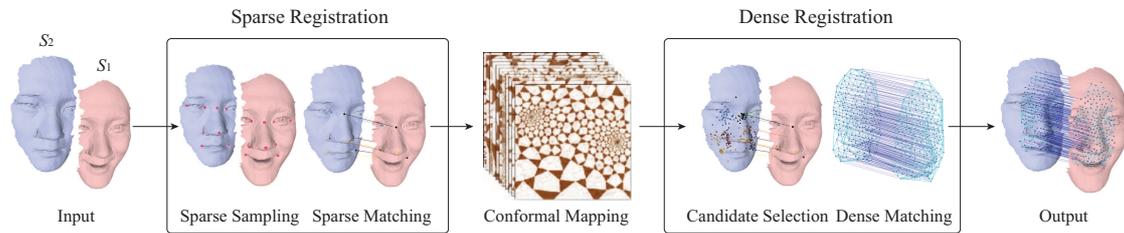


Figure 1: Overview of our algorithmic framework for surface registration as described in Sec. 1.2.

search for the correspondences using probabilistic formulations based on geodesic distances. Nevertheless, issues like the computational complexity and inaccuracy of geodesic distance towards establishing dense correspondences reduce the applicability of these methods to a limited number of correspondences. An alternate approach for searching correspondences in the intrinsic space is to parameterize the shape into a high dimensional space such that its Euclidean distance approximates the intrinsic properties of the surface. For example, the idea of diffusion map based methods [51, 58] is to represent the shape in an embedding space that preserves the commute time, which has the advantage of being robust to topological changes. To overcome the high computational complexity in establishing point correspondences, the idea of functional maps [50] was proposed to establish the correspondences in the functional space, which is only limited to shapes with bijective mapping.

Graph matching is a powerful framework for establishing feature correspondences [3, 18], able to combine multiple matching metric (*e.g.*, appearance similarity and geometric compatibility) through the integration of singleton, pairwise or higher-order interactions between nodes into a potential function (*e.g.*, [24, 67]). It has been employed by many existing works (*e.g.*, [6, 7]) to encode contextual constraints between feature points for high-level applications (mostly for images) such as object recognition. However, its use in 3D surface matching is very limited, probably due to the fact that more sophisticated metrics are required in order to model the spatial information between points on the surfaces [69]. Numerous optimization methods have been proposed to solve the graph matching problem, such as spectral relaxation methods (*e.g.*, [21, 46]), methods based on continuous relaxation (*e.g.*, [56, 66]), *etc.* Besides, [67] proposed a novel pairwise graph-matching algorithm based on the dual-decomposition framework, which can provide certain optimality guarantee on the solution.

Markov Random Fields (MRFs) [70], in particular pairwise MRFs, have been widely applied in many fields in computer vision, such as image segmentation, stereo, detection and registration, *etc.* [9]. Recently, the applications of higher-order MRFs in computer vision have become more and more popular, along with the development of inference methods (*e.g.*, [32, 43, 75]). Compared with pairwise models, higher-order MRFs allow a better characterization of statistics among random variables (*e.g.*, [36, 54]), *e.g.*, when we want to model measurements involving more than two variables (*e.g.*, [4, 28]) and/or to characterize invariant statistics [71, 72].

1.2 System Overview

Let us now provide a brief overview (Fig. 1) of the proposed surface registration system that handles the dense correspondence problem in a hierarchical manner: sparse surface registration and dense surface registration. Preliminary versions of this system were described in our earlier work [78, 79].

In the sparse feature matching stage, an initial set of sparse feature points are selected without loss of generality according to geometric principles (among the local maxima of Gaussian curvature [47] or the average geodesic distance function [35] from the input surfaces S_1 and S_2). On these points we find

n_s correspondences between the two feature sets $p_i^1 \leftrightarrow p_i^2, i = 1, \dots, n_s$, using a higher-order graph matching algorithm (Sec. 2).

Once such correspondences have been established, the dense feature matching stage relies on them to constrain the local search space for each point on the surface (Sec. 3). Since every three correspondences determine a unique conformal map between two surfaces, n_s sparse correspondences give us $\binom{n_s}{3}$ candidate matching points on S_2 for any point $p \in S_1$. These points are then pruned and clustered to obtain meaningful matching candidates for each point (Sec. 4.1.2). Given the discrete set of candidate correspondences for each point on S_1 , the dense surface registration problem becomes combinatorial. To impose constraints on the deformation field, we introduce a generic local surface deformation model defined on the triangulated graph of these points on S_1 with meaningful candidate matching points (Sec. 3.1). Finally, a higher-order MRF optimization is then formulated and solved to obtain the optimal dense registration result (Sec. 3.2 and Sec. 3.3).

2 Sparse surface registration using higher-order graph matching

In this section, we present our sparse surface registration algorithm based on a higher-order graph matching formulation. First of all, we introduce the higher-order graph matching problem and in particular a general pseudo-boolean formulation.

2.1 Pseudo-boolean higher-order graph matching

Let us denote by P_1 and P_2 two sets of points, and $P \equiv P_1 \times P_2$ the set of possible correspondences between P_1 and P_2 . We define the boolean indicator variable

$$x_a = \begin{cases} 1 & \text{if } a = (p_a^1, p_a^2) \in P \text{ is an active correspondence,} \\ 0 & \text{otherwise,} \end{cases}$$

where $p_a^1 \in P_1$ and $p_a^2 \in P_2$ denote the two points involved in correspondence a . A basic constraint imposed on the matching configuration is that each point in P_1 is mapped to at most one point in P_2 , while for each point in P_2 there is at most one point in P_1 mapping to it. Therefore, we have the set of feasible solutions defined as,

$$\mathcal{C} = \{\mathbf{x} \in \{0, 1\}^{P_1 \times P_2} \mid \sum_{i \in P_1} x_{i,j} \leq 1, \sum_{j \in P_2} x_{i,j} \leq 1, \forall i \in P_1 \text{ and } \forall j \in P_2\}. \quad (1)$$

Note that a point is allowed to have no active correspondence in order to deal with partial matching. With the above constraints, the higher-order graph matching problem can be formulated as follows:

$$\min_{\mathbf{x} \in \mathcal{C}} \{E(\mathbf{x}|\theta) = \sum_{a \in P} \theta_a x_a + \sum_{(a,b) \in P \times P} \theta_{ab} x_a x_b + \sum_{(a,b,c) \in P \times P \times P} \theta_{abc} x_a x_b x_c\}, \quad (2)$$

where θ_a is the matching cost for each correspondence $a \in P$, θ_{ab} for a pair of correspondences $(a, b) \in P \times P$, and θ_{abc} for a triplet of correspondences $(a, b, c) \in P \times P \times P$. In fact, the matching constraints can be reduced to pairwise terms in the energy function. More specifically, we observe the following equivalence:

$$\begin{aligned} & \forall i \in P_1, \sum_{j \in P_2} x_{i,j} \leq 1 \\ & \text{iff } \min_{x_{i,j}} \sum_{j', j'' \in P_2, j' \neq j''} \theta^{\infty} x_{i,j'} x_{i,j''} = 0 \end{aligned} \quad (3)$$

where θ^∞ is a sufficiently large number. We use $P^C = \{(a, b) | a, b \in P, a \neq b \text{ and } (p_a^1 = p_b^1 \text{ or } p_a^2 = p_b^2)\}$ to denote the set of pairs that encodes the matching constraints for all the correspondences. Thus, the general higher-order matching problem can be formulated as the following pseudo-boolean optimization problem [10]:

$$\min_{\mathbf{x} \in \{0,1\}^{P_1 \times P_2}} \{E(\mathbf{x}|\theta) = \sum_{a \in P} \theta_a x_a + \sum_{(a,b) \in P \times P} \theta_{ab} x_a x_b + \sum_{(a,b) \in P^C} \theta^\infty x_a x_b + \sum_{(a,b,c) \in P \times P \times P} \theta_{abc} x_a x_b x_c\}. \quad (4)$$

The above formulation is general and is able to handle partial matching by properly defining the potentials.

Because of the positive weight θ^∞ that encodes the matching constraint, the energy function is non-convex[11], and in general this is an NP-hard problem [10]. The advantage of the pseudo-boolean formulation is that theoretically any higher-order terms can be reduced into a quadratic term [10] which can be done efficiently [32]. For its optimization, we adopt the flexible dual-decomposition technique [42]. Before discussing the details of the optimization technique, we first introduce the application of a higher-order graph matching formulation to the problem of surface matching.

2.2 Higher-order graph matching for sparse surface matching

In the context of the above general higher-order graph matching formulation, the singleton terms define both appearance and geometric similarities, the pairwise terms constrain the matching solution space (see Eq. 1) and the higher-order terms encode the intrinsic deformation errors.

Singleton potentials For each correspondence (p_a^1, p_a^2) we consider both the geometric and texture information to define its potential as in [65]. We consider the Gaussian curvature $\text{curv}(p)$ at point p , which is invariant under isometric transformations [22], and the texture value $\text{tex}(p)$ at point p if available. The singleton potential for a correspondence (p_a^1, p_a^2) is defined as

$$\theta_i = (\text{curv}(p_a^1) - \text{curv}(p_a^2))^2 + \lambda_0 (\text{tex}(p_a^1) - \text{tex}(p_a^2))^2, \quad (5)$$

where λ_0 balances the weight between the curvature and the texture information. Similarly, other geometric features can also be considered such as spin-images [33].

Pairwise potentials We use pairwise potentials to encode the mapping constraints for the graph matching (see Sec. 2.1), by setting θ^∞ and θ_{ab} in Eq. 4 to be 10^5 and 0, respectively, in our experiments.

Higher-order potentials The uniformization theorem [25] states any 3D surface can be flattened conformally to a canonical 2D domain. Within such a mapping each feature point p has a parametric coordinate in the complex plane $z_p \in \hat{\mathbb{C}}$. The flexibility of this conformal mapping is represented by a Möbius transformation, which can be uniquely determined by fixing the mappings of any three points on the surface to the 2D parametrization domain. Inspired by [47], we compute the matching score between two triplets as the deformation error based on the Möbius transform as detailed below.

Given two surfaces, S_1 and S_2 , and a mapping from $\{p_a^1, p_b^1, p_c^1\} \subset S_1$ to $\{p_a^2, p_b^2, p_c^2\} \subset S_2$, we first determine the associated Möbius transformation $m^1(z)$ and $m^2(z)$ that map each triplet to a prefixed configuration $\{e^{i\frac{2\pi}{3}}, e^{i\frac{4\pi}{3}}, e^{i2\pi}\} \subset \hat{\mathbb{C}}$, where $\hat{\mathbb{C}}$ represents the complex domain. Such transformations essentially equip each point on the surface S_1 and S_2 with a new coordinate in $\hat{\mathbb{C}}$. Let us denote the new coordinate for each point p as $z(p) \in \hat{\mathbb{C}}$.

Similar to [47], we establish correspondences between the two sets $P_1 \subset S_1$ and $P_2 \subset S_2$ by searching for the mutually closest point correspondence set under the new coordinates, denoted as:

$$\begin{aligned} \mathcal{M}_{abc} = & \{(p_1, p_2) | p_1 \in P_1, p_2 \in P_2, \\ & \forall p'_2 \in P_2 \setminus \{p_2\}, |z(p_1) - z(p_2)| < |z(p_1) - z(p'_2)|, \\ & \forall p'_1 \in P_1 \setminus \{p_1\}, |z(p_1) - z(p_2)| < |z(p'_1) - z(p_2)|\} \end{aligned}$$

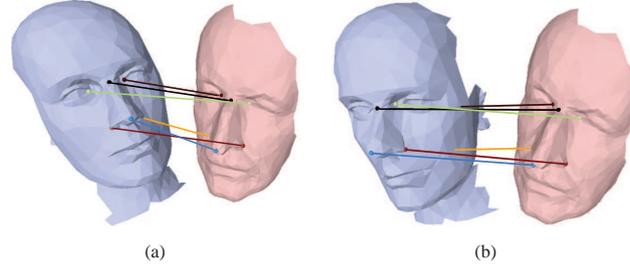


Figure 2: An example shows the ambiguity by considering only the intrinsic information. The matching scores in (a) and (b) are the same from Eq. 6 based on the Möbius transform, since the distances between the matching features are identical. However, such ambiguity can be avoided by adding the extrinsic similarity information (e.g., normal and curvature).

and define the deformation error as

$$E_{abc} = \sum_{(p_1, p_2) \in \mathcal{M}_{abc}} |z(p_1) - z(p_2)|^2. \quad (6)$$

Note that both \mathcal{M}_{abc} and E_{abc} are computed by aligning p_k^1 to p_k^2 , where $k \in \{a, b, c\}$. Given \mathcal{M}_{abc} and E_{abc} , we then define the Möbius matching potential as follows,

$$\theta_{abc}^{\text{Möbius}} = \begin{cases} \frac{E_{abc}}{|\mathcal{M}_{abc}|^2} - 1 & \text{if } \frac{E_{abc}}{|\mathcal{M}_{abc}|} < \delta \\ 1/|\mathcal{M}_{abc}| & \text{otherwise} \end{cases}. \quad (7)$$

Here $|\mathcal{M}_{abc}|$ denotes the number of valid correspondences and δ is a threshold to rule out unlikely correspondences (in our experiment $\delta = 0.1$), without which the minimization problem of Eq. 2 would encourage as many correspondences as possible even when some of them do not match. Intuitively, the more matching pairs and the smaller the distance between those mutually closest pairs, the lower the potential energy.

However, considering the Möbius energy alone can introduce certain ambiguity since it only encodes isometric information (an example is shown in Fig. 2). To resolve such an ambiguity, we consider the Gaussian map of the surface. The Gaussian map is defined as the mapping of the normal at each point on the surface to the unit sphere [22], which encodes the extrinsic geometric information of the surface. In order to avoid ambiguities in orientation, the order of each triplet on the Gaussian map is constrained. Two triplets have the same orientation if and only if the determinant of their normals have the same sign. Therefore, we have

$$\theta_{abc}^{\text{Gaussian}} = \begin{cases} 0 & \text{if } \det(\mathbf{n}_a^1, \mathbf{n}_b^1, \mathbf{n}_c^1) \det(\mathbf{n}_a^2, \mathbf{n}_b^2, \mathbf{n}_c^2) \geq 0 \\ 1/|\mathcal{M}_{abc}| & \text{otherwise} \end{cases}.$$

Here $\mathbf{n}_i \in \mathbb{R}^3$ denotes the normal at point i , and $\det(\mathbf{n}_a, \mathbf{n}_b, \mathbf{n}_c)$ denotes the determinant of the 3×3 matrix $[\mathbf{n}_a, \mathbf{n}_b, \mathbf{n}_c]$. This is introduced as a soft constraint in our framework, because in the extreme case, the normal can reverse its orientations when the surface undergoes very large deformations. Finally, the triple potential for each possible triple matching $(p_a^1, p_b^1, p_c^1) \rightarrow (p_a^2, p_b^2, p_c^2)$ can be defined as

$$\theta_{abc} = \lambda_1 \theta_{abc}^{\text{Möbius}} + \lambda_2 \theta_{abc}^{\text{Gaussian}}. \quad (8)$$

After the potentials of the graph matching problem (Eq. 4) are defined, we next discuss its optimization.

2.3 Optimization and computational complexity

The idea of dual-decomposition is to re-formulate the original problem as the union of several sub-problems that are easier to solve [41, 42]. For the graph matching problem in Eq. 2, let θ denote the vector of the weights of the singleton, pairwise and triplet terms, and I denote the set of subproblems. The decomposition is represented by $E(\mathbf{x}|\theta) = \sum_{\sigma \in I} \rho_{\sigma} E^{\sigma}(\mathbf{x}|\theta^{\sigma})$, where ρ_{σ} is the weight for each subproblem. Then the original problem is solved by updating the parameter θ^{σ} of each subproblem σ that increases the energy of the dual problem. Moreover, we have the following decomposition constraint:

$$\sum_{\sigma \in I} \rho_{\sigma} \theta^{\sigma} = \theta. \quad (9)$$

If we can find a lower bound $\Phi_{\sigma}(\theta^{\sigma})$ for each subproblem, *i.e.*, $\Phi_{\sigma}(\theta^{\sigma}) \leq \min_{\mathbf{x}} E^{\sigma}(\mathbf{x}|\theta^{\sigma})$, then we can obtain a lower bound for the original problem, *i.e.*,

$$\Phi(\theta) = \sum_{\sigma \in I} \rho_{\sigma} \Phi_{\sigma}(\theta^{\sigma}) \leq \sum_{\sigma \in I} \rho_{\sigma} E^{\sigma}(\mathbf{x}|\theta^{\sigma}) = E(\mathbf{x}|\theta). \quad (10)$$

This lower bound is to be maximized using projected subgradient method so that a solution to the original problem can be extracted from the Lagrangian solutions [42].

In particular, we decompose the original problem into the following three subproblems:

1. a **linear subproblem** which considers only the singleton term $\sum_{a \in P} \theta_a x_a$. This linear subproblem is also known as the linear assignment problem and can be solved efficiently by Hungarian algorithm [3].
2. a **higher-order pseudo-boolean subproblem** by reducing the higher-order terms in 2 to quadratic terms [10] which can be solved by the QPBO algorithm [39]. Here we employ [32] for the reduction.
3. a **local subproblem** which divides the original surface into small regions and uses an exhaustive search to find the optimal solution in each small region.

Opposite to [67] that only considers pairwise subproblems, a higher-order pseudo-boolean subproblem is introduced because of the higher-order terms in Eq. 2. After solving the subproblems, the dual variables $\{\theta^{\sigma}\}$ are updated by projecting to the space that satisfies Eq. 9 [42, 67].

For the dual-decomposition algorithm above, the most expensive step in each iteration is the max-flow computation. In our paper, we employ the popular implementation in [44], whose worst case complexity is $O(mn^2|C|)$, where m is the number of edges, n is the number of vertices, and $|C|$ is the cost of minimum cut. Assuming we select $|P_1|$ and $|P_2|$ feature points from two surfaces, there are $O(|P_1|^3|P_2|^3)$ possible triplets, each represented by a higher-order term in Eq. 2.

3 Dense surface registration using higher-order MRFs

Given the solution to the sparse surface matching problem presented in the previous section, we now look at the problem of dense surface registration. The main obstacles in extending the sparse matching approach to dense come from (i) *Model accuracy*: the optimality properties of the obtained solution degrade as the sampling points goes dense, since the Möbius matching potential defined in Eq. 7 is only discriminating when the features are far apart. A new potential function is desired that is able to encode the localized, anisometric deformations in the local search space obtained from the sparse matching. (ii) *Computational complexity*: when the number of sampling points n becomes large, it is computationally expensive to solve the higher-order graph matching problem since the graph structure would grow in the order of $O(n^6)$ if we consider all possible triplets.

To overcome the above issues, we build up a novel higher-order MRF-based dense surface registration method by introducing a set of new elements: (i) a new deformation model; (ii) a higher-order MRF model for determining the optimal instance of this deformation model; (iii) an efficient matching candidate selection scheme; (iv) an efficient optimization algorithm for the MRF inference. In the rest of this section, we present these elements successively.

3.1 A generic deformation model

We will first introduce a mathematical formulation that accurately characterizes arbitrary surface deformations for a dense set of points on the surface. We start with our generic deformation model in the continuous setting and then derive its meaning in the discrete setting.

3.1.1 Continuous setting

Let $(\mathcal{M}, g_{\mathcal{M}})$ denote a surface \mathcal{M} equipped with a Riemannian metric $g_{\mathcal{M}}$. In Riemannian geometry [23], a surface is defined by its local charts $\mathcal{M} = U_{\alpha} \cup U_{\beta} \dots$, and each open subset U_{α} is in one to one correspondences $\phi_{\alpha} : U_{\alpha} \rightarrow \mathbb{R}^2$, where ϕ_{α} is the local *parametrization*. For any $p \in U_{\alpha} \subset \mathcal{M}$, a metric tensor is associated with p as a symmetric positive definite matrix:

$$g^{\alpha}(p) = \begin{pmatrix} g_{11}^{\alpha}(p) & g_{12}^{\alpha}(p) \\ g_{21}^{\alpha}(p) & g_{22}^{\alpha}(p) \end{pmatrix}. \quad (11)$$

In order for different local representations to describe the same surface, the following chain rule must be satisfied:

$$\forall p \in U_{\alpha} \cap U_{\beta}, g^{\alpha}(p) = J_{\alpha\beta}(p)^T g^{\beta}(q) J_{\alpha\beta}(p). \quad (12)$$

Here $J_{\alpha\beta}$ is the Jacobian matrix of the transformation between the local coordinate systems of U_{α} and U_{β} . Any local representation satisfying this transformation rule is a valid parametrization of the surface. Since the metric tensor at any point $p \in \mathcal{M}$ is positive definite, we can always apply a proper linear transformation to its parametrization ϕ_{α} such that $g^{\alpha}(p)$ is the identity matrix. Such a parametrization is called the *canonical parametrization* for p :

Definition 1. (Canonical parametrization) For any $p \in \mathcal{M}$, a parametrization $\phi_{\alpha} : U_{\alpha} \rightarrow \mathbb{R}^2$ is called canonical parametrization for p if the metric tensor at p is the identity matrix.

Accordingly, the Jacobian matrix J_{pq} between the two points p and q under their canonical parametrizations is called the *canonical Jacobian*. In the following, we will show that considering the canonical parametrization/Jacobian allows us to characterize arbitrary deformations between surfaces that is independent from both the intrinsic and extrinsic representations of the surface, which is the main advantage of our deformation model.

Now we consider arbitrary diffeomorphisms between the parametrization domains of two surfaces. For any correspondence $p \in U_{\alpha} \subset \mathcal{M} \rightarrow q \in U_{\beta} \subset \mathcal{N}$, the change of metric $g^{\alpha}(p) \rightarrow J_{\alpha\beta}(p)^T g^{\beta}(q) J_{\alpha\beta}(p)$ reflects how an infinitesimal circle is deformed into an infinitesimal ellipse. In particular, under canonical parametrizations for points p and q (i.e., both $g^{\alpha}(p)$ and $g^{\beta}(q)$ are identity matrices), the matrix $J_{pq}^T J_{pq}$ accurately characterizes such local deformation, where J_{pq} is the canonical Jacobian between p and q . If we only consider the change of shape, i.e., how a circle is deformed into an ellipse regardless of its orientation, the distortion along its two principle directions can be represented by the two eigenvalues λ_1 and λ_2 of $J_{pq}^T J_{pq}$ (Fig. 3 (a)). Therefore, the local deformation between two surfaces can be characterized by such two eigenvalues λ_1, λ_2 for each pair (p, q) of corresponding points. Formally, we define:

Definition 2. (Canonical distortion coefficients) The canonical distortion coefficients (CDCs) between two point p and q are defined as the eigenvalues of the Jacobian transformation matrix $J_{pq}^T J_{pq}$ between any canonical parametrization at p and q .

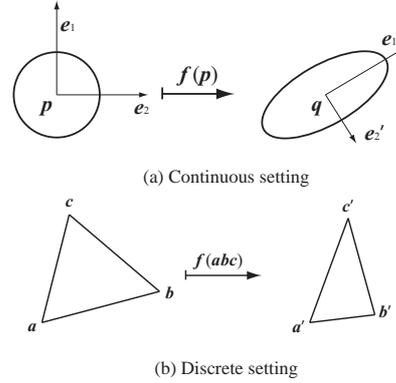


Figure 3: The finite element method assumes the transformation between facets to be piecewise linear and $f(\vec{ab}) = \vec{a'b'}$, $f(\vec{ac}) = \vec{a'c'}$. Under the linearity assumption, the Jacobian can be computed in a closed form for each pair of triangular facets $\triangle abc \mapsto \triangle a'b'c'$.

CDCs are generic deformation features that are able to characterize a wide class of deformation groups. For instance, below are two typical classes of deformations that can be characterized by CDCs:

1. In the case of the isometric deformation, a unit circle is mapped to a unit circle, *i.e.*, $\lambda_1 = \lambda_2 = 1$.
2. In the case of the conformal deformation, a unit circle can be mapped to a circle with arbitrary radius [61]. Thus, $\lambda_1 = \lambda_2 \neq 0$.

To further connect the canonical distortion coefficient to a general class of diffeomorphisms defined in the complex plane $f : \mathcal{U}_\alpha \rightarrow \mathcal{U}_\beta$, between any canonical parametrization x_α and x_β for p and q respectively, we define

$$\begin{aligned} \frac{\partial f}{\partial z} &= \frac{1}{2} \left(\frac{\partial x_\alpha^1}{\partial x_\beta^1} + \frac{\partial x_\alpha^2}{\partial x_\beta^2} \right) + \frac{i}{2} \left(\frac{\partial x_\alpha^2}{\partial x_\beta^1} - \frac{\partial x_\alpha^1}{\partial x_\beta^2} \right) \\ \frac{\partial f}{\partial \bar{z}} &= \frac{1}{2} \left(\frac{\partial x_\alpha^1}{\partial x_\beta^1} - \frac{\partial x_\alpha^2}{\partial x_\beta^2} \right) + \frac{i}{2} \left(\frac{\partial x_\alpha^2}{\partial x_\beta^1} + \frac{\partial x_\alpha^1}{\partial x_\beta^2} \right), \end{aligned}$$

where we assume $x_\alpha = x_\alpha^1 + ix_\alpha^2$ and $x_\beta = x_\beta^1 + ix_\beta^2$. The notion of quasi-conformality [2] can be characterized by the *Beltrami-coefficient*:

$$\mu(z) \equiv \frac{\partial f}{\partial \bar{z}} / \frac{\partial f}{\partial z},$$

which gives us all the information about the conformality of f . Without loss of generality, let us suppose $\lambda_1 \geq \lambda_2$. It can be shown that $|\mu(z)| = (\sqrt{\lambda_1} - \sqrt{\lambda_2}) / (\sqrt{\lambda_1} + \sqrt{\lambda_2})$. In particular, f is called *holomorphic* if $\mu(z) = 0$ ([27]), *i.e.*, $\lambda_1 = \lambda_2$, coinciding with the fact that a holomorphic function is another description of conformal mapping. Hence, the Beltrami-coefficient generalizes conformal mapping and can be partially determined using CDCs. However, the Beltrami-coefficient is for surface parametrization, where the scaling factor is lost. The proposed CDC preserves the scale information which is important for shape matching. Besides, unlike the Beltrami-coefficient, the CDC is directly extendable to nD .

3.1.2 Discrete setting through finite elements

The basic assumption in finite element analysis [12] is that the continuous space can be approximated using a set of *basis elements* (*e.g.*, polynomial functions defined on each facet) with continuity preserved

at the boundaries among the basis elements. We consider the most common representation of a continuous surface – a triangular mesh, with triangular facets as basis finite elements. In this discrete setting, CDCs are assumed to be constant for each basis element (*i.e.*, each triangular facet). Thus, the concept of canonical parametrization can be expressed in the following manner: a parametrization of a point p is locally Euclidean at p if the images of any two tangent vectors have the same angle and length. In the discrete setting, this means in the canonical parametrization domain, a 3D surface facet $\triangle abc$ is mapped onto 2D by preserving all the angles and edge lengths.

Regarding the computation of the canonical Jacobian, in the continuous setting, the Jacobian matrix at a point p is a linear transformation that maps between the tangent spaces at p and q , respectively. Given a basis element $\triangle abc$ in the discrete setting, the tangent space at p corresponds to the linear space spanned by $\triangle abc$. Hence, the linear mapping $J(\cdot)$ between two canonical domains $\triangle abc$ and $\triangle a'b'c'$ should satisfy $J(\vec{ab}) = a'\vec{b}'$ and $J(\vec{ac}) = a'\vec{c}'$, which can be computed in closed-form. Since $J(\cdot)$ is linear, $J(\vec{bc}) = b'\vec{c}'$ must be hold true, *i.e.*, the Jacobian for mapping $p \rightarrow q$ in the continuous case corresponds to a linear transformation mapping $\vec{ab} \rightarrow a'\vec{b}'$, $\vec{ac} \rightarrow a'\vec{c}'$ in the discrete case (Fig. 3).

Alg. 1 summarizes the algorithm for computing CDCs. When the shape is n -manifold, the computation of CDCs only requires solving n linear equations and eigenvalues. Note that although the computation looks analogous to [53] and [55] for surface parametrization due to the piecewise linear assumption, Alg. 1 is derived in the context of Riemannian geometry for shape deformation.

Algorithm 1: Algorithm for computing the canonical distortion coefficients (CDCs) for each triangular facet.

Input : $\triangle abc$ and its mapping $\triangle a'b'c'$

Output : CDCs for mapping from $\triangle abc$ to $\triangle a'b'c'$.

Step One: Map the triangles $\triangle abc$ and $\triangle a'b'c'$ to 2D while keeping their orientation.

Step Two: Compute the 2×2 linear transformation J mapping \vec{ab} to $a'\vec{b}'$ and \vec{ac} to $a'\vec{c}'$.

Step Three: Compute the eigenvalues, λ_1 and λ_2 of $J^T J$.

Step Four: Output λ_1 and λ_2

Once the representation and deformation model have been defined, we can now introduce a generic MRF formulation for dense shape registration.

3.2 MRF formulation for shape registration

The choice of representation (CDCs) allows us to either deform the original shape (*e.g.*, [60]), or determine the correspondences between two shapes, S_1 and S_2 , by combining with other matching cues. Assuming that a triangulated set of n points $\mathcal{V} = \{p_u | p_u \in S_1, u = 1, \dots, n\}$ are sampled on the surface S_1 , the task of shape registration becomes determining the correspondence for each point $p \in \mathcal{V}$ on S_2 , which we solve via a higher-order MRF framework.

The structure of our higher-order MRF model can be represented as a hyper-graph $\mathcal{G} = (\mathcal{V}, \mathcal{F})$, corresponding to the triangulation of the set of points on the surface S_1 , where \mathcal{V} denotes the vertex set and $\mathcal{F} \subset \mathcal{V} \times \mathcal{V} \times \mathcal{V}$ denotes the triangular facet set. We also associate a random variable x_u for each vertex $u \in \mathcal{V}$ to represent the set of matching candidates of the vertex u on the shape S_2 . A configuration/matching for the point set \mathcal{V} can therefore be denoted by $\mathbf{x} = (x_u)_{u \in \mathcal{V}}$. For the sake of clarity and simplicity, x_u will also be used to denote the corresponding point on S_2 .

Regarding the MRF energy, we first define the unary potential function $\theta_u(x_u)$ as the difference in the feature descriptor (*e.g.*, texture or shape context) between u and its correspondence x_u :

$$\theta_u(x_u) = |\text{fea}_{S_1}(u) - \text{fea}_{S_2}(x_u)|^2, \quad (13)$$

where $\text{fea}_{\mathcal{S}}(\cdot) \in \mathbb{R}^n$ denotes the feature vector at a point on shape \mathcal{S} .

Next, let $\lambda_{uvw}(x_u, x_v, x_w) \in \mathbb{R}^2$ denote the CDCs computed from deforming Δ_{uvw} to $\Delta_{x_u x_v x_w}$ (Alg. 1), the higher-order potential can then be defined as follows:

$$\theta_{uvw}(x_u, x_v, x_w) = \rho(\lambda_{uvw}(x_u, x_v, x_w)), \quad (14)$$

where $\rho(\cdot)$ encodes the deformation constraints on the CDC values (detailed discussion is given in Sec. 3.2.1). Given the above potential functions, shape registration becomes searching for the optimal configuration \mathbf{x} that minimizes

$$E(\mathbf{x}) = \sum_{u \in \mathcal{V}} \theta_u(x_u) + \sum_{(u,v,w) \in \mathcal{F}} \theta_{uvw}(x_u, x_v, x_w). \quad (15)$$

In the following, we first discuss the practical aspects of imposing the deformation constraints (Eq. 14), before we talk about its optimization.

3.2.1 Imposing deformation constraints

To constrain the deformation of a surface represented by CDCs, we assume that *the deformations are similar across different shapes of the same type*. Such an assumption has been successfully applied in shape [62] and facial expressions [68] transfer. The ground truth deformation prior can be obtained by 3D scanning systems with reliable texture information (e.g., markers). Fig. 4 shows an example of human facial expressions. The two 3D data in (a) and (c) are captured with markers using the system introduced in [74]. Here we select two frames with the largest expression difference to measure the maximal possible change of CDCs. Fig. 4 (b), (d), (f) and (g) show the visualization of the distribution of CDCs. From the above data set we obtain the allowed bound for human face expression changing from neutral to large deformation as $I_1 = [0.7, 5.66]$, $I_2 = [0.1, 4]$ for λ_1 and λ_2 , respectively. For the problem of surface registration, we define a Potts-like energy for the higher-order terms in Eq. 15 as follows:

$$\theta_{uvw}(x_u, x_v, x_w) = \begin{cases} 0 & \text{if } \lambda_1 \in I_1 \text{ and } \lambda_2 \in I_2 \\ 10 & \text{otherwise} \end{cases}, \quad (16)$$

where λ_1 and λ_2 are the CDCs obtained by matching Δ_{uvw} to $\Delta_{x_u x_v x_w}$.

3.2.2 Matching candidate set

Now we discuss the definition of $x_u, u \in \mathcal{V}$ in Eq. 15. A candidate selection process is necessary since including all the points on S_2 to x_u would be computationally prohibitive. A common workaround is to prune off unlikely matching points based on neighborhood information, as shown in the 2D graph matching case [24]. However, for the 3D surface matching problem, a surface may undergo very large deformation. Therefore neighborhood cannot be straightforwardly defined based on their 3D coordinates. To overcome this, we observe that *once three correspondences are known, an alignment of the two surfaces in the conformal mapping domain can be obtained* in a closed form by determining the associated Möbius transformation. Especially, the alignment is accurate when the two surfaces undergo perfect isometric deformation and the error increases continuously as the deformation deviates from isometry. Hence, it is natural to get the matching candidate set for solving the dense surface registration based on the sparse registration results (Fig. 1). The implementation details are discussed in Sec. 4.1.2.

3.3 An efficient higher-order MRF optimization

Efficient inference in higher-order MRFs is a very active research problem and various techniques have been proposed. Most of existing methods are based on *order-reduction* [32, 26] (first reduce higher-order terms in to pairwise ones and then solve the problem using *graph cuts* techniques [39]) or *dual-decomposition* [41]. However, the algorithms designed for general MRFs often lack efficiency in terms of

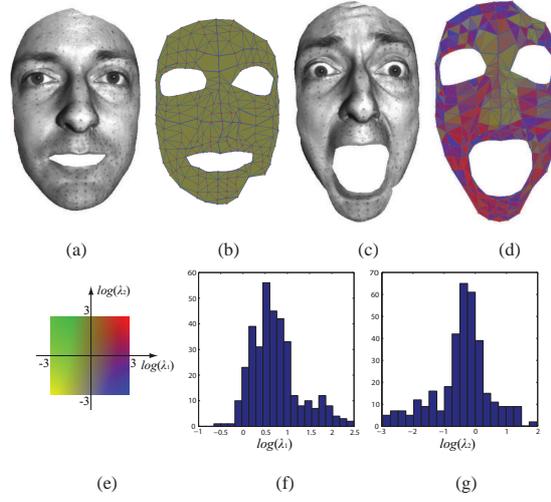


Figure 4: Expression deformation prior is obtained by 3D scanned data with markers. (a) and (c) show the 3D scan of the onset and peak of a facial expression with large shape deformations respectively. (b) and (d) are the corresponding triangular templates constructed from the 3D scan data. (d) shows the deformation from (b) to (d) represented by CDCs and illustrated by a color map shown in (e). The histogram of the CDC values are shown in (f) and (g).

computation and/or memory when solving MRFs with special structure and/or potential energy functions. In order to efficiently perform the inference in our MRF model (Sec. 3.2), we explore the structural property and develop a parallel optimization algorithm that requires minimal memory and achieves significant speedup via GPU¹.

Let us first look at the dual problem for the LP relaxation of the minimization problem of the energy in Eq. 15, by introducing an indicator variable $\tau_{u;i}$ for $u \in \mathcal{V}$, $i \in \mathcal{L}$:

$$\tau_{u;i} = \begin{cases} 1 & \text{if } x_u = i \\ 0 & \text{otherwise} \end{cases},$$

Similarly, $\tau_{uvw;ijk}$ is introduced for each $(u, v, w) \in \mathcal{F}$ and $(i, j, k) \in \mathcal{L} \times \mathcal{L} \times \mathcal{L}$:

$$\tau_{uvw;ijk} = \begin{cases} 1 & \text{if } x_u = i, x_v = j, x_w = k \\ 0 & \text{otherwise} \end{cases}.$$

By defining $\theta_{u;i} = \theta_u(i)$ and $\theta_{uvw;ijk} = \theta_{uvw}(i, j, k)$, we obtain the following integer LP formulation for the minimization problem of the energy in Eq. 15:

$$\begin{aligned} \min_{\tau} \quad & \sum_{u \in \mathcal{V}} \sum_{i \in \mathcal{L}} \theta_{u;i} \tau_{u;i} + \sum_{(u,v,w) \in \mathcal{F}} \sum_{(i,j,k) \in \mathcal{L}^3} \theta_{uvw;ijk} \tau_{uvw;ijk} \\ \text{s.t.} \quad & \sum_i \tau_{u;i} = 1, & \forall u \in \mathcal{V} \\ & \sum_{i,j,k} \tau_{uvw;ijk} = 1, & \forall (u, v, w) \in \mathcal{F} \\ & \sum_{j,k} \tau_{uvw;ijk} = \tau_{u;i}, & \forall (u, v, w) \in \mathcal{F} \text{ and } i \in \mathcal{L} \\ & \tau_{u;i}, \tau_{uvw;ijk} \in \{0, 1\}. \end{aligned}$$

¹Note that the proposed method could also be combined with existing methods for improving message passing in pairwise MRFs (e.g., [45, 40]) to boost further the efficiency, which is under investigation.

By relaxing the domains of the variables $\tau_{u;i}$ and $\tau_{uvw;ijk}$ to $[0, 1]$, we obtain the dual form of the LP problem as follows:

$$\begin{aligned} \max_M \sum_u \min_i \bar{\theta}_{u;i} + \sum_{(u,v,w) \in \mathcal{F}} \min_{i,j,k} \bar{\theta}_{uvw;ijk} & \quad (17) \\ \text{s.t. } \bar{\theta}_{u;i} = \theta_{u;i} + \sum_{(u,v,w) \in \mathcal{F}} M_{uvw;u;i}, \forall u \in \mathcal{V} \text{ and } i \in \mathcal{L} \\ \bar{\theta}_{uvw;ijk} = \theta_{uvw;ijk} - M_{uvw;u;i} - M_{uvw;v;j} - M_{uvw;w;k}, \\ \forall (u,v,w) \in \mathcal{F} \text{ and } (i,j,k) \in \mathcal{L} \times \mathcal{L} \times \mathcal{L}. \end{aligned}$$

Here $M_{uvw;u;i}$ is the dual variable (message) corresponding to the constraint $\sum_{j,k} \tau_{uvw;ijk} = \tau_{u;i}$ (Fig. 7(a)).

The dual problem of Eq. 17 can be solved by *min-sum diffusion* algorithm [75] as shown in Alg. 2. It has been shown that at convergence, the solution satisfies *J-consistency* condition as introduced in [75]. Since after each update of the message, only a reparameterization of the MRF is performed, no extra memory is needed for storing all the dual variables $M_{uvw;u;i}$. Hence, the memory requirement for the Alg. 3 is only for storing primal variables, *i.e.*, $O(|V||\mathcal{L}| + |\mathcal{F}||\mathcal{L}|^3)$.

Algorithm 2: Min-sum diffusion algorithm.

```

repeat
  for each  $M_{uvw;u;i}$  do
     $M_{uvw;u;i} = \frac{1}{2}[\theta_{u;i} - \min_{j,k} \theta_{uvw;ijk}]$  and reparameterize  $\theta_{u;i}$  and  $\theta_{uvw;ijk}$  according to the
    constraints in Eq. 17.
  end for
until convergence

```

Each update of the message in Alg. 2 only involves the parameters in a triangle. Also within each facet Δuvw , the update of each label $M_{uvw;u;i}$, $i = \{1, \dots, L\}$ is independent. Hence the algorithm can be significantly accelerated.

In order to explore the parallelism of the min-sum algorithm 2, we define the concept of independent face set.

Definition 3. (Independent face set) Given a graph $\mathcal{G} = (\mathcal{V}, \mathcal{F})$, a subset $\mathcal{F}_k \subset \mathcal{F}$ is called independent face set if for any $f_i, f_j \in \mathcal{F}_k$, $f_i \cap f_j = \emptyset$.

The decomposition of a set \mathcal{F} into subsets of independent face sets $\mathcal{F} = \cup_i \mathcal{F}_i$ can be efficiently computed in polynomial time by a simple greedy algorithm. Hence we can implement Alg. 2 in parallel as shown in Alg. 3. The maximal speedup achieved in Alg. 3 is $\max_i (|\mathcal{F}_i|/|\mathcal{L}|)$.

Algorithm 3: Parallel min-sum diffusion algorithm.

```

Decompose  $\mathcal{F}$  into independent face sets  $\cup_i \mathcal{F}_i$ 
repeat
  for each Independent face set  $\mathcal{F}_i$ , in parallel for all  $(u,v,w) \in \mathcal{F}_i$  and  $k \in \mathcal{L}$  do
    Update the message  $M_{uvw;u;k}$ ,  $M_{uvw;v;k}$  and  $M_{uvw;w;k}$  and do reparameterization (Alg. 2).
  end for
until convergence

```

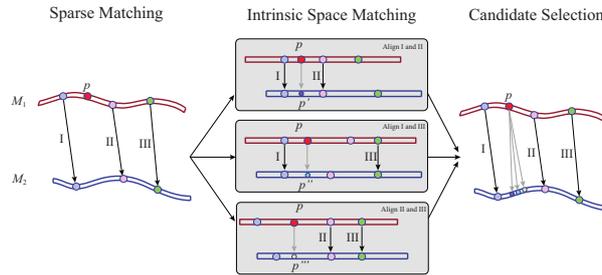


Figure 5: 2D Illustration of candidate selection scheme.

4 Experimental results

4.1 Implementation

Our system is tested on an Intel[®] Xeon(TM) 3.4G PC with 4G RAM and an NVIDIA[®] Geforce 9800GTX+ graphics card. Let us first discuss the technical aspects of the solution prior of introducing the obtained results.

4.1.1 Sparse graph matching

For the potential functions of the graph matching algorithm defined in Sec. 2.2, the weights of Eq. 5 and 8 are defined as $\lambda_0 = 1$, $\lambda_1 = 0.1$ and $\lambda_2 = 1$.

The mid-edge uniformization algorithm [47, 53] is used for the conformal mapping. The computation of mid-edge uniformization involves solving a symmetric linear equation, which can be efficiently computed by GPU [17] (it takes $< 1s$ for a mesh with 10^4 faces). For each input mesh, only one conformal mapping is computed using the uniformization algorithm [47, 53]. Any other conformal mappings are computed in a closed form by the Möbius transformation induced by three-correspondence matches (Sec. 2.2). For the higher-order graph matching algorithm (Sec. 2.3), the convergence of the dual-decomposition optimization depends on the input features. In our experiments, we observe that the more outliers (un-matched points), the more iterations it takes to converge.

4.1.2 Candidate selection & clustering

Given the sparse correspondences between S_1 and S_2 , the goal of candidate selection is to obtain matching candidates on S_2 for each point on S_1 . For every triplet of sparse correspondences, we map them to a prefixed configuration by solving a Möbius transformation, using the same procedure described in Sec. 2.2. By aligning the triplet correspondences in the 2D domain, an alignment of the two surfaces is also obtained, giving us a matching point on S_2 for every point on S_1 . This way, n sparse matching correspondences produce $\binom{n_s}{3}$ candidate matching points for every point on S_1 . A 2D illustration of the candidate selection scheme is given in Fig. 5.

Our candidate selection approach is robust to sparse matching errors, since only part of the sparse matching results are used for selecting each dense matching candidate. A qualitative evaluation is shown in Fig. 6. We can see that most of the candidate points are close to the desired correspondences on the target surface. It should be noted that considering all triplet correspondences provides an overly complete set of matching candidates. We further reduce the search space by clustering these dense matching candidates using the mean shift clustering [20] to find the modes of this density. Compared to parametric representations, mean shift does not require nonlinear optimization to learn the distribution parameters.

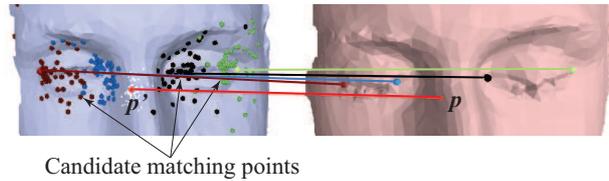


Figure 6: An example showing candidate points obtained from different Möbius transforms and their clustering. For any point p from the source surface, the clustering of its matching candidate points on the target surface gives us a matching candidate p' .

Given the sparse matching result, the candidate selection and clustering of 10^3 points based on 10 sparse features takes around 1 minute. The input to the dense surface registration stage are those vertices \mathcal{V} on the source surface with at least one valid matching point from our candidate selection stage and a triangulation of these vertices is constructed based on their parametrization, resulting in a higher-order graph $\mathcal{G} = (\mathcal{V}, \mathcal{F})$. In our experiment, the aforementioned candidate selection process typically provides 2 – 4 candidate points for each $v \in \mathcal{V}$. To guarantee matching accuracy, we uniformly re-sample $L = 64$ points for each $v \in \mathcal{V}$ near the original matching candidates. The computation of all the L^3 possible CDCs for one facet takes only 2.0ms on average using GPU. The computation of the potential $\theta_{uvw;ijk}$ for a graph with $|\mathcal{F}| = 2000$ takes 3s.

4.1.3 MRF optimization of the higher-order dense graph matching

We implement Alg. 3 using the NVIDIA[®] CUDA architecture [34]. In approximation algorithms, the approximation error (AE) is defined as the gap between the optimal integral solution and the solution obtained by the algorithm. In order to test the AE, we design the test inputs as follows: Given any input mesh, we randomly assign a ground truth label l_u for each node $u \in \mathcal{V}$. We define the singleton potentials of Eq. 15 as

$$\theta_u(x_u) = \begin{cases} 0 & \text{if } x_u = l_u \\ rnd(1) & \text{otherwise} \end{cases},$$

where $rnd(1)$ is a random number between $[0, 1]$. Also we define the higher-order potentials as

$$\theta_{uvw}(x_u, x_v, x_w) = \begin{cases} 0 & \text{if } (x_u, x_v, x_w) = (l_u, l_v, l_w) \\ rnd(1) & \text{otherwise} \end{cases}.$$

In this case, the optimal solution of Eq. 15 should be $(l_u)_{u \in \mathcal{V}}$. Fig. 7 (a) shows the result of our algorithm using the above designed test cases for different mesh and label size. Although the total energy increases with mesh size, the average energy per term (vertex and facet) remains significantly low (< 0.01 for all cases). Fig. 7 (b) shows the comparison on average time taken per iteration, between the implementations with and without GPU accelerations. The total number of iterations depends on the shape of the energy potential. In our experiment, the algorithm converges within 3000 iterations. Hence our algorithm is both memory and computationally efficient.

4.2 Experimental validation

4.2.1 Sparse & dense shape registration

In this section, we evaluate our new algorithmic framework for dense surface matching. Regarding the optimization, in our experiments, we match surfaces with large deformations and inconsistent boundaries

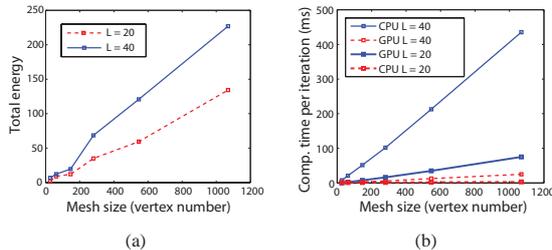


Figure 7: Performance analysis of our MRF optimization algorithm. (a) shows the optimality using the test cases described in Sec. 4.1.3. (b) shows the speedup using the parallel implementation of Alg. 3. L is the number of labels for each node. We show the runtime per iteration since different inputs have different iteration counts.

(partial overlapping). The number of vertices for each mesh is in the range of 1,500 – 4,000. With our higher-order graph based system, we can find the dense matching for 60 – 90 percent of all vertices, which is illustrated as matched/total (no. of matched vertices/no. of total vertices of the source surface) for each example. The lion data of Fig. 9 comes from [62] and the face and hand data are captured with texture by the 3D scanner introduced in [74]. To measure the quality of dense registration, from the Delaunay triangulation of the points on the source surface, we consider the ratio of the area of each local triangle to the area of its matched triangle. For the natural deformations (*e.g.*, expression change, stretched arms or bending figures) we experimented with, the local area is not expected to undergo abrupt change. Therefore the area ratio is expected to be close to one for every local triangle.

Matching with largely inconsistent boundaries and partial overlapping: The mid-edge uniformization algorithm allows to map the boundaries of the surface to slits and preserve the conformal structure of the surface in an exact sense. Hence it is suitable for matching partially overlapping surfaces. This property can be combined with our candidate selection scheme to determine the outliers near the boundary where the mean shift clustering returns a low score. Examples are shown in Fig. 10, 11, and 12. An example of significant non-overlap between the two meshes is shown in Fig. 1.

Matching with large deformations: Fig. 8, 9, 12 and 13 show results matching two surfaces undergoing a large deformation. Even when the sparse features can not all be selected consistently (as shown in Fig. 13), our higher-order graph matching algorithm in Sec. 2.3 is able to find reliable sparse correspondences (Fig. 13(a)) and obtain a dense surface matching result through the two-stage optimization scheme in Sec. 1.2 (Fig. 13(b)).

Comparison experiments: Fig. 11 shows a comparison between our algorithm and the Least Square Conformal Mapping (LSCM) approach [73]. Although LSCM can handle free boundaries, there is no theoretical guarantee that the conformal structure is preserved near the boundary and it might include self-intersections in the mapping [59]. In our comparison, we use the feature correspondences computed from the sparse matching stage to initialize the LSCM experiments. The inaccuracy of the LSCM approach can be observed in Fig. 11(c). In this example, although all vertices on the left mesh are matched to the right mesh, there are approximately 42 percent flipped triangles. Note that here we cannot compare directly with the results in [73] where the initial feature points were manually selected.

Furthermore, we compare our results with the recent dense surface matching approach [35] whose code is publicly available online. Similar to [35], we define the matching error for a set of test points \mathcal{P}_{test} to be

$$\text{Err}(f, f_{\text{true}}) = \sum_{p \in \mathcal{P}_{test}} d_{S_2}(f(p), f_{\text{true}}(p)) / |\mathcal{P}_{test}| \quad (18)$$

where $d_{S_2}(\cdot, \cdot)$ denotes the geodesic distance between two points on surface S_2 normalized by $\sqrt{\text{Area}(S_2)}$. Here we obtain the “ground truth” map f_{true} by manually selecting the matching points for around 100

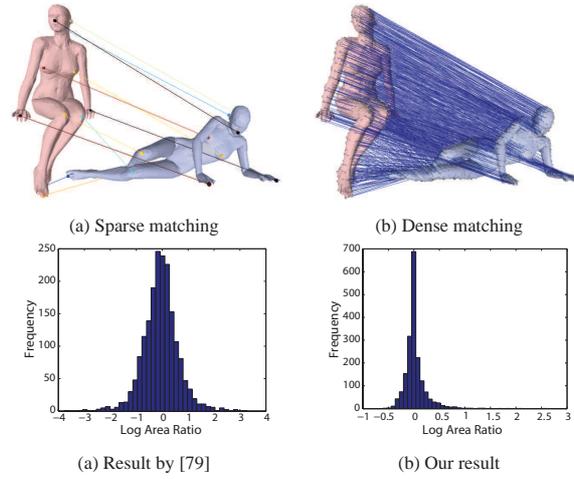


Figure 8: Matching result for the body data: (matched/total = 2861/3376)

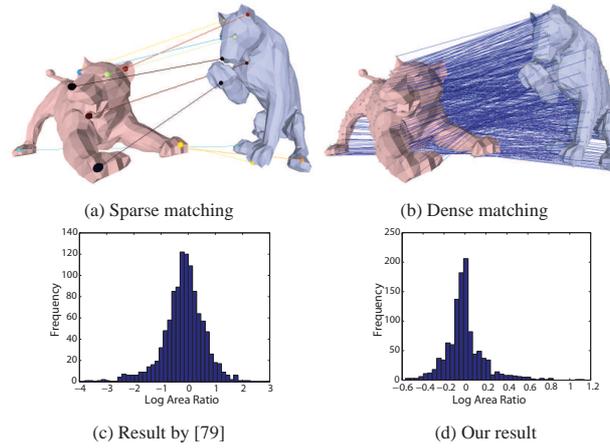


Figure 9: Matching result for the lion data: (matched/total = 1105/1251)

points on S_1 . To reduce the error caused by individual bias, we average the matching results by 5 people for each point. The result of the comparison between our method and [35] is shown in Table 1. Note that [35] assumes the mapping between two surfaces to be bijective and there is no explicit underlying deformation model in selecting the final dense match. In contrast, we take into account the partial matching problem both in our graph-based formulation and candidate selection scheme, and an accurate deformation model is considered in selecting the optimal dense matching result.

4.2.2 Sparse & dense surface tracking

We have also applied our approach to the tracking of dynamic, 3D scanned data. For the singleton term in Eq. 15, we employed the robust metric proposed in [80]. Both the consistencies between consecutive frames and between current frame and the first frames are taken into account. To impose inter-frame consistency, we use the same data set as [80] and select the two consecutive frames with the largest deformation change to obtain the range of CDCs between frames, *i.e.*, $I_1 = [0.874, 1.143]$ and $I_2 =$

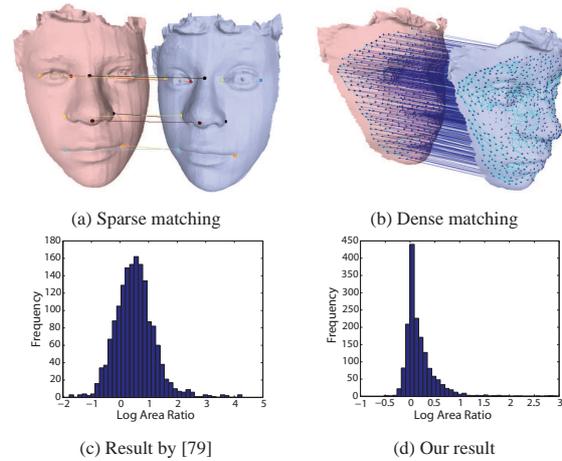


Figure 10: Matching result for the face data: (matched/total = 2098/2644)

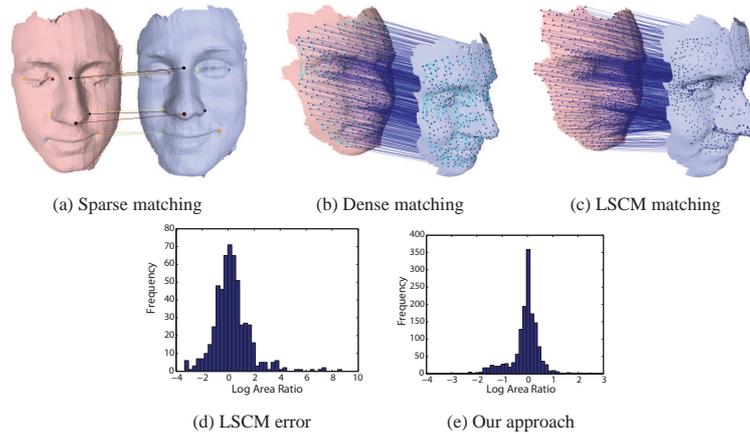


Figure 11: Comparison with LSCM approach [73] for dense surface matching. (matched/total = 1455/1635) (best viewed in color). Notice the high number of flipped triangles in (c)

$[0.846, 1.182]$ for λ_1 and λ_2 respectively (Eq. 16). Also we handle drift errors by imposing consistency between the first frame and the current frame, using the same deformation prior learned in Fig. 4.

Fig. 14 shows the tracking results on the BU-4DFE database [76], which consists of 3D dynamic facial expressions of different subjects. A mesh template is constructed in the first frame and tracked in the subsequent frames. Because of the temporal continuity in consecutive frames, sufficient matching candidates can be obtained by only looking at the neighborhood of each point. In this dataset, the texture information is noisy and thus only relying on texture information can easily lead to erroneous results. Nevertheless, with our deformation model encoded in the higher-order terms of the MRF model, we have achieved excellent tracking results for sequences with significant anisometric facial deformation as shown in Fig. 14.

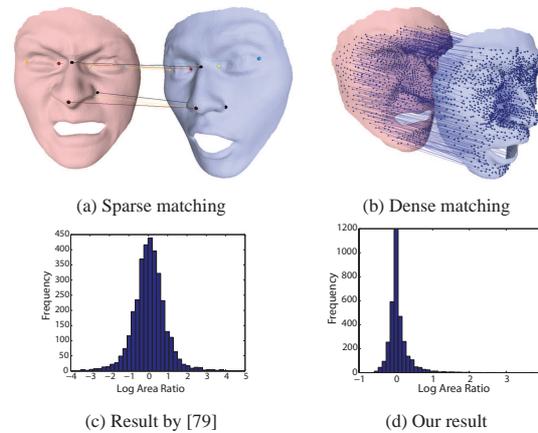


Figure 12: Dense matching under large non-rigid deformations. (matched/total = 2378/3633)

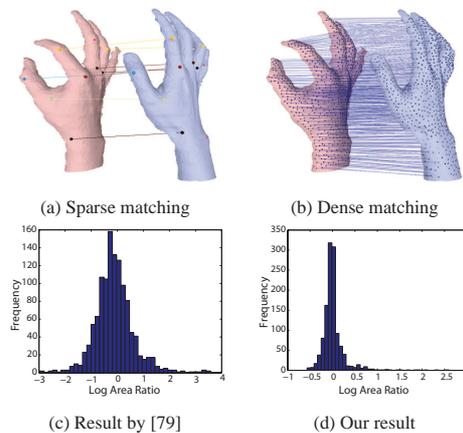


Figure 13: Dense matching under multiple articulated deformations. (matched/total = 1224/1786)

5 Conclusion

We proposed a higher-order graph-based approach for dense, non-rigid surface registration. Specifically, a two-stage algorithm is proposed to constrain the search space and improve matching accuracy, through an efficient candidate selection scheme and an accurate deformation model respectively. In our sparse surface registration stage, a higher-order graph matching formulation is proposed to combine the appearance and geometry similarities as well as the intrinsic deformation error based on isometry assumption. In our dense surface registration stage, a generic deformation model, encoded by a higher-order MRF formulation, is proposed to take into account detailed, anisometric surface deformations. Meanwhile, we developed two efficient algorithms for the higher-order graph matching and MRF problems, respectively. As a result, our proposed method achieves robust dense registration between non-rigid surfaces with large deformations, partial overlapping and inconsistent boundaries and scale.

This work provides a solid basis for multiple future directions. For example, our matching method can be applied to dynamic 3D shape completion used in 3D virtual video conference, thank to its ability to partially match surfaces with large deformations. Our deformation model, namely CDCs, can be used for driving the animation of 3D objects, by solving the 3D embedding of a template object using

Data	Our method	Kim <i>et.al.</i> [35]
Body (Fig. 8)	0.0622	0.2437
Lion (Fig. 9)	0.0832	0.1790
Face (Fig. 10)	0.0319	0.0465
Face (Fig. 12)	0.0565	0.0865
Hand (Fig. 13)	0.0481	0.1193

Table 1: Comparison with a latest intrinsic space based approach for dense surface registration [35]. The average error is calculated based on Eq. 18.

our higher-order MRF optimization. We will also explore more sophisticated ways to learn the CDCs and apply them in the problems of expression transfer, face recognition, *etc.* Meanwhile, we are also particularly interested in exploring uncertainly in our higher-order graph-matching framework based on existing techniques (*e.g.*, [37, 48]), towards further enhancing the power of such a framework and boosting further its performance.

Acknowledgments

This work was partially supported from the European Research Council Starting Grant DIOCLES (ERC-STG-259112), AFOSR FA9550-10-1-0294, NSF DMS-1221339, and NSF DMS-1418255.

References

- [1] Microsoft[®] Kinect, 2010.
- [2] L. V. Ahlfors. *Lectures on Quasiconformal Mappings*. American Mathematical Society, 2 edition, 2006.
- [3] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin. *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall, 1993.
- [4] A. A. Amini, T. E. Weymouth, and R. C. Jain. Using dynamic programming for solving variational problems in vision. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 12(9):855–867, 1990.
- [5] D. Anguelov, P. Srinivasan, H.-C. Pang, D. Koller, S. Thrun, and J. Davis. The correlated correspondence algorithm for unsupervised registration of nonrigid surfaces. In *Proc. Neural Information Processing Systems*, 2004.
- [6] S. Belongie, J. Malik, and J. Puzicha. Shape matching and object recognition using shape contexts. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 24(4):509–522, Apr. 2002.
- [7] A. C. Berg, T. L. Berg, and J. Malik. Shape matching and object recognition using low distortion correspondences. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pages 26–33, 2005.
- [8] P. J. Besl and N. D. McKay. A method for registration of 3-D shapes. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 14(2):239–256, 1992.
- [9] A. Blake, P. Kohli, and C. Rother. *Markov Random Fields for Vision and Image Processing*. MIT Press, 2011.

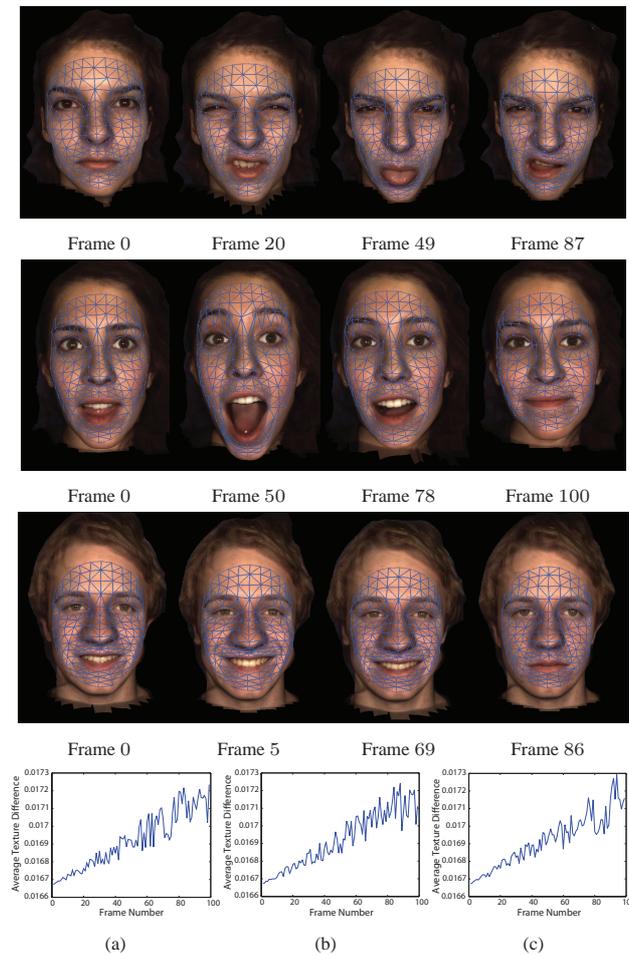


Figure 14: Shape tracking results. The average texture difference for all correspondences between every frame and the first frame for the three sequences, from top to bottom are shown in (a), (b) and (c), respectively.

- [10] E. Boros and P. L. Hammer. Pseudo-boolean optimization. *Discrete Applied Mathematics*, 123(1-3):155–225, 2002.
- [11] S. P. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- [12] S. C. Brenner and R. Scott. *The Mathematical Theory of Finite Element Methods*. Springer, 3 edition, 2007.
- [13] A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *Proc. National Academy of Sciences*, 103:1168–1172, 2006.
- [14] A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Expression-invariant representations of faces. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 2004:1042–1053, 2007.
- [15] A. M. Bronstein, M. M. Bronstein, and R. Kimmel. *Numerical Geometry of Non-Rigid Shapes*. Springer, 2008.

- [16] B. J. Brown and S. Rusinkiewicz. Global non-rigid alignment of 3-D scans. *ACM Trans. Graph.*, 26, 2007.
- [17] L. Buatois, G. Caumon, and B. Levy. Concurrent number cruncher: a GPU implementation of a general sparse linear solver. *Int. J. Parallel Emerg. Distrib. Syst.*, 24(3):205–223, 2009.
- [18] R. Burkard, M. Dell’Amico, and S. Martello. *Assignment Problems*. Society for Industrial and Applied Mathematics, 2009.
- [19] R. J. Campbell and P. J. Flynn. A survey of free-form object representation and recognition techniques. *Computer Vision and Image Understanding*, 81(2):166–210, 2001.
- [20] D. Comaniciu and P. Meer. Mean shift: A robust approach toward feature space analysis. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 24(5):603–619, 2002.
- [21] T. Cour, P. Srinivasan, and J. Shi. Balanced graph matching. In *Proc. Neural Information Processing Systems*, 2007.
- [22] M. P. do Carmo. *Differential Geometry of Curves and Surfaces*. Prentice Hall, 1976.
- [23] M. P. do Carmo. *Riemannian Geometry*. Birkhäuser, 1992.
- [24] O. Duchenne, F. Bach, I. Kweon, and J. Ponce. A tensor-based algorithm for high-order graph matching. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2009.
- [25] H. M. Farkas and I. Kra. *Riemann Surfaces*. Springer, 2004.
- [26] A. Fix, J. Chen, E. Boros, and R. Zabih. Approximate MRF inference using bounded treewidth subgraphs. In *Proc. European Conf. Computer Vision*, volume 7572, pages 385–398, 2012.
- [27] T. Frankel. *The Geometry of Physics—An Introduction*. Cambridge, 2 edition, 2004.
- [28] B. Glocker, T. H. Heibel, N. Navab, P. Kohli, and C. Rother. TriangleFlow: Optical Flow with Triangulation-Based Higher-Order Likelihoods. In *Proc. European Conf. on Computer Vision*, 2010.
- [29] D. Hahnel, S. Thrun, and W. Burgard. An extension of the ICP algorithm for modeling nonrigid objects with mobile robots. In *Proc. of the Int’l Joint Conference on Artificial Intelligence*, pages 915–920, 2003.
- [30] Q.-X. Huang, B. Adams, M. Wicke, and L. J. Guibas. Non-rigid registration under isometric deformations. *Symposium on Geometry Processing*, pages 1449–1457, 2008.
- [31] X. Huang, N. Paragios, and D. N. Metaxas. Shape registration in implicit spaces using information theory and free form deformations. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 28(8):1303 – 1318, 2006.
- [32] H. Ishikawa. Higher-order clique reduction in binary graph cut. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2009.
- [33] A. Johnson. *Spin-Images: A Representation for 3-D Surface Matching*. PhD thesis, Carnegie Mellon University, 1997.
- [34] A. Keane. CUDA (compute unified device architecture). 2006.
- [35] V. G. Kim, Y. Lipman, and T. Funkhouser. Blended intrinsic maps. *ACM Trans. Graph. (Proc. of SIGGRAPH 2011)*, 30:79:1–79:12, 2011.

- [36] P. Kohli, M. Pawan Kumar, and P. H. S. Torr. P3 & beyond: move making algorithms for solving higher order functions. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 31(9):1645–1656, Sept. 2009.
- [37] P. Kohli and P. H. S. Torr. Measuring uncertainty in graph cut solutions. *Computer Vision and Image Understanding*, 112(1):30–38, 2008.
- [38] D. Koller and N. Friedman. *Probabilistic Graphical Models: Principles and Techniques*. The MIT Press, 2009.
- [39] V. Kolmogorov and C. Rother. Minimizing nonsubmodular functions with graph cuts—a review. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 29(7):1274–1279, 2007.
- [40] N. Komodakis. Towards more efficient and effective lp-based algorithms for MRF optimization. In *Proc. European Conf. Computer Vision*, pages 520–534, 2010.
- [41] N. Komodakis and N. Paragios. Beyond pairwise energies: Efficient optimization for higher-order MRFs. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2009.
- [42] N. Komodakis, N. Paragios, and G. Tziritas. MRF optimization via dual decomposition: Message-passing revisited. In *Proc. IEEE Int’l Conf. Computer Vision*, 2007.
- [43] N. Komodakis, N. Paragios, and G. Tziritas. MRF energy minimization and beyond via dual decomposition. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 33(3):531–552, Mar. 2011.
- [44] N. Komodakis, G. Tziritas, and N. Paragios. Performance vs computational efficiency for optimizing single and dynamic MRFs: Setting the state of the art with primal-dual strategies. *Computer Vision and Image Understanding*, 112(1):14–29, 2008.
- [45] M. P. Kumar and P. H. S. Torr. Fast memory-efficient generalized belief propagation. In *Proc. European Conf. Computer Vision*, pages 451–463, 2006.
- [46] M. Leordeanu and M. Hebert. A spectral technique for correspondence problems using pairwise constraints. In *Proc. IEEE Int’l Conf. Computer Vision*, pages 1482–1489, 2005.
- [47] Y. Lipman and T. Funkhouser. Möbius voting for surface correspondence. *ACM Trans. Graph.*, 28(3):72:1–72:12, 2009.
- [48] K. Miller, M. P. Kumar, B. Packer, D. Goodman, and D. Koller. Max-margin min-entropy models. In *Proc. Int’l Conf. Artificial Intelligence and Statistics*, pages 779–787, 2012.
- [49] H. E. A. E. Munim, A. A. Farag, and A. A. Farag. Shape representation and registration in vector implicit spaces: Adopting a closed form solution in the optimization process. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 35(3):763 – 768, 2012.
- [50] M. Ovsjanikov, M. Ben-Chen, J. Solomon, A. Butscher, and L. Guibas. Functional maps: a flexible representation of maps between shapes. *ACM Trans. Graph.*, 31(4):30:1–30:11, 2012.
- [51] M. Ovsjanikov, Q. Merigot, F. Memoli, and L. J. Guibas. One point isometric matching with the heat kernel. *Comput. Graph. Forum*, 29(5):1555–1564, 2010.
- [52] N. Paragios, M. Rousson, and V. Ramesh. Non-rigid registration using distance functions. *Computer Vision and Image Understanding*, 89(203):142–165, 2003.
- [53] U. Pinkall and K. Polthier. Computing discrete minimal surfaces and their conjugates. *Experimental Mathematics*, 2(1):15–36, 1993.

- [54] S. Roth and M. J. Black. Fields of Experts. *Int'l J. Computer Vision*, 82(2):205–229, Jan. 2009.
- [55] P. V. Sander, J. Snyder, S. J. Gortler, and H. Hoppe. Texture mapping progressive meshes. In *SIGGRAPH '01*, pages 409–416, 2001.
- [56] C. Schellewald and C. Schnorr. Probabilistic subgraph matching based on convex relaxation. In *Energy Minimization Methods in Computer Vision and Pattern Recognition (EMMCVPR)*, 2005.
- [57] A. Shaji, A. Varol, L. Torresani, and P. Fua. Simultaneous point matching and 3D deformable surface reconstruction. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2010.
- [58] A. Sharma, R. P. Horaud, J. Cech, and E. Boyer. Topologically-robust 3D shape matching based on diffusion geometry and seed growing. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2011.
- [59] A. Sheffer, E. Praun, and K. Rose. Mesh parameterization methods and their applications. *Found. Trends. Comput. Graph. Vis.*, 2(2):105–171, 2006.
- [60] O. Sorkine and M. Alexa. As-rigid-as-possible surface modeling. In *Symposium on Geometry Processing*, pages 109–116, 2007.
- [61] K. Stephenson. *Introduction to Circle Packing: The Theory of Discrete Analytic Functions*. Cambridge University Press, 2005.
- [62] R. W. Sumner and J. Popović. Deformation transfer for triangle meshes. In *SIGGRAPH*, pages 399–405, 2004.
- [63] A. Tevs, A. Berner, M. Wand, I. Ihrke, and H.-P. Seidel. Intrinsic shape matching by planned landmark sampling. In *Eurographics*, volume 30, pages 543–552, 2011.
- [64] A. Tevs, M. Bokeloh, M. Wand, A. Schilling, and H.-P. Seidel. Isometric registration of ambiguous and partial data. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pages 1185 – 1192, 2009.
- [65] N. Thorstensen and R. Keriven. Non-rigid shape matching using geometry and photometry. In *Proc. Asian Conf. Computer Vision*, 2009.
- [66] P. H. S. Torr. Solving Markov Random Fields using semi definite programming. In *Ninth Int'l Workshop on Artificial Intelligence and Statistics*, 2003.
- [67] L. Torresani, V. Kolmogorov, and C. Rother. Feature correspondence via graph matching: Models and global optimization. In *Proc. European Conf. on Computer Vision*, volume 5303, pages 596–609. Springer, 2008.
- [68] D. Vlasic, M. Brand, H. Pfister, and J. Popović. Face transfer with multilinear models. *ACM Trans. Graph.*, 24:426–433, 2005.
- [69] C. Wang, M. M. Bronstein, A. M. Bronstein, and N. Paragios. Discrete minimum distortion correspondence problems for non-rigid shape matching. In *International Conference on Scale Space and Variational Methods in Computer Vision (SSVM)*, 2011.
- [70] C. Wang, N. Komodakis, and N. Paragios. Markov Random Field modeling, inference & learning in computer vision & image understanding: A survey. *Computer Vision and Image Understanding*, 117(11):1610–1627, 2013.

- [71] C. Wang, O. Teboul, F. Michel, S. Essafi, and N. Paragios. 3D knowledge-based segmentation using pose-invariant higher-order graphs. In *International Conference, Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2010.
- [72] C. Wang, Y. Zeng, L. Simon, I. Kakadiaris, D. Samaras, and N. Paragios. Viewpoint invariant 3D landmark model inference from monocular 2D images using higher-order priors. In *Proc. IEEE Int'l Conf. Computer Vision*, 2011.
- [73] S. Wang, Y. Wang, M. Jin, X. D. Gu, and D. Samaras. Conformal geometry and its applications on 3D shape matching, recognition, and stitching. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 29(7):1209–1220, 2007.
- [74] Y. Wang, M. Gupta, S. Zhang, S. Wang, X. Gu, D. Samaras, and P. Huang. High resolution tracking of non-rigid 3D motion of densely sampled data using harmonic maps. In *Proc. IEEE Int'l Conf. Computer Vision*, 2005.
- [75] T. Werner. Revisiting the linear programming relaxation approach to gibbs energy minimization and weighted constraint satisfaction. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 32:1474–1488, 2010.
- [76] L. Yin, X. Chen, Y. Sun, T. Worm, and M. Reale. A high-resolution 3D dynamic facial expression database. In *Face and Gesture Recognition*, 2008.
- [77] W. Zeng, Y. Zeng, Y. Wang, X. Yin, X. Gu, and D. Samaras. 3D non-rigid surface matching and registration based on holomorphic differentials. In *Proc. European Conf. on Computer Vision*, 2008.
- [78] Y. Zeng, C. Wang, X. Gu, D. Samaras, and N. Paragios. A generic deformation model for dense non-rigid surface registration: a higher-order mrf-based approach. In *Proc. IEEE Int'l Conf. Computer Vision*, pages 3360 – 3367, 2013.
- [79] Y. Zeng, C. Wang, Y. Wang, X. Gu, D. Samaras, and N. Paragios. Dense non-rigid surface registration using high-order graph matching. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pages 382–389, 2010.
- [80] Y. Zeng, C. Wang, Y. Wang, X. Gu, D. Samaras, and N. Paragios. Intrinsic dense 3D surface tracking. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2011.
- [81] D. Zhang and M. Hebert. Harmonic maps and their applications in surface matching. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 1999.
- [82] L. Zhang, N. Snavely, B. Curless, and S. M. Seitz. Spacetime faces: high resolution capture for modeling and animation. *ACM Trans. Graph.*, 23(3):548–558, 2004.

Contents

1	Introduction	3
1.1	Related work	4
1.2	System Overview	5
2	Sparse surface registration using higher-order graph matching	6
2.1	Pseudo-boolean higher-order graph matching	6
2.2	Higher-order graph matching for sparse surface matching	7
2.3	Optimization and computational complexity	9
3	Dense surface registration using higher-order MRFs	9
3.1	A generic deformation model	10
3.1.1	Continuous setting	10
3.1.2	Discrete setting through finite elements	11
3.2	MRF formulation for shape registration	12
3.2.1	Imposing deformation constraints	13
3.2.2	Matching candidate set	13
3.3	An efficient higher-order MRF optimization	13
4	Experimental results	16
4.1	Implementation	16
4.1.1	Sparse graph matching	16
4.1.2	Candidate selection & clustering	16
4.1.3	MRF optimization of the higher-order dense graph matching	17
4.2	Experimental validation	17
4.2.1	Sparse & dense shape registration	17
4.2.2	Sparse & dense surface tracking	19
5	Conclusion	21



**RESEARCH CENTRE
SACLAY – ÎLE-DE-FRANCE**

Parc Orsay Université
4 rue Jacques Monod
91893 Orsay Cedex

Publisher
Inria
Domaine de Voluceau - Rocquencourt
BP 105 - 78153 Le Chesnay Cedex
inria.fr

ISSN 0249-6399