



HAL
open science

Higher-order Graph Principles towards Non-rigid Surface Registration

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

► **To cite this version:**

Yun Zeng, Chaohui Wang, Xianfeng Gu, Dimitris Samaras, Nikos Paragios. Higher-order Graph Principles towards Non-rigid Surface Registration. [Research Report] RR-8607, INRIA Saclay - Ile-de-France; INRIA. 2015, pp.31. hal-01086052v2

HAL Id: hal-01086052

<https://inria.hal.science/hal-01086052v2>

Submitted on 30 Sep 2015

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

ISRN INRIA/RR--8607--FR+ENG

ISSN 0249-6399



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 — version 2 — initial version September 2014 —
revised version September 2015 — 28 pages

Abstract: This report casts surface registration as the problem of finding a set of discrete correspondences through the minimization of an energy function, which is composed of geometric and appearance matching costs, as well as higher-order deformation priors. Two higher-order graph-based formulations are proposed under different deformation assumptions. The first formulation encodes isometric deformations using 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 sparse feature point sets on surfaces undergoing highly anisometric deformations. Nevertheless, its performance degrades significantly when addressing anisometric registration for a set of densely sampled points. This issue is rigorously addressed subsequently through a novel deformation model that is able to handle arbitrary diffeomorphisms between two surfaces. Such a deformation model is introduced into a higher-order Markov Random Field for dense surface registration, and is inferred using 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. A series of experiments demonstrate the accuracy and the efficiency of the proposed framework, notably in challenging cases of large and/or anisometric deformations, or surfaces that are partially occluded.

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

* Google Inc.

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

‡ Department of Computer Science, Stony Brook University, Stony Brook, NY, 11794, USA

§ 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, qui est composée de fonctions de coûts mesurant 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 nouveau 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, et être inféré en utilisant 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, Appariement de Graphes d'ordre élevé, Géométrie Conforme, Champs de Markov Aléatoires d'ordre élevé

1 Introduction

Finding the correspondences between two or more surfaces is a prerequisite for many applications in computer vision, medical imaging and computer graphics such as shape recognition, deformation transfer, object recognition and segmentation [18]. Furthermore, the proliferation of 3D content (*e.g.*, [1, 49, 65]) has further enhanced the need of developing a robust surface registration method, in particular for noisy-sampled 3D shapes undergoing large and/or non-rigid deformations.

Surface registration approaches can be classified as either *extrinsic* (*i.e.*, the coordinates for each point on the surface is known, *e.g.*, [8]), or *intrinsic* (*i.e.*, the surface is represented in a parametrization space, *e.g.*, [15]). Registration of surfaces undergoing large deformations in the extrinsic space is intractable due to the large number of degrees of freedom present in their non-rigid deformations. In contrast, intrinsic 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 common case, intrinsic methods assume that two surfaces undergo *isometric* (*i.e.*, distance or metric preserving) deformations. Such an assumption is a reasonable approximation for most of the real-world deformations. State-of-the-art intrinsic methods refer to geodesic/exponential maps (*e.g.*, [13, 14]), conformal maps (*e.g.*, [64, 68, 72]) and diffusion maps (*e.g.*, [44, 50]). In particular, conformal maps provide a closed-form solution to the dense surface matching problem, and can be directly generalized to anisometric deformations using quasi-conformal maps [2], which can handle arbitrary diffeomorphisms between two surfaces. Despite the fact that conformal maps are well-suited for establishing dense correspondences between surfaces undergoing large deformations, the use of a single map is unreliable in the presence of inconsistent boundary conditions, anisometric deformation, and noise. It is therefore necessary to consider correspondences obtained from multiple mappings and use them towards robust metrics able to handle uncertainties in the input. Graph-based approaches have strong advantages in handling noise and occlusions [9, 61]. and thus will be used as the basis for our approach.

In this paper, we cast the surface registration problem as a two-stage matching task that seeks correspondences between two discrete point sets, sampled from two surfaces. In such a setting, deformation constraints between two surfaces are represented by up to third-order interactions among these point sets. Thus, the optimal registration corresponds to the optimal solution of an objective function defined over a higher-order graph. To efficiently couple the strength of intrinsic surface registration methods and graph-based approaches, we propose two graph-based formulations to solve sparse and dense surface registration, leading to both computational efficiency and accurate dense surface registration results.

Initially, surface registration is expressed as a higher-order graph matching problem that integrates both extrinsic matching costs (*e.g.*, texture, curvature and normal consistencies) and intrinsic deformation constraints (deviation from isometry, inspired by [42]), and is solved using the dual-decomposition technique [36]. Despite its success in registering sparsely sampled points, the deformation metric [42] is only robust to points on the surface that are sufficiently sparse. It is no longer discriminating enough, when the set of points to be registered becomes dense.

To this end, we propose a novel local surface deformation model to characterize arbitrary deformations between two surfaces. According to Riemannian geometry [23], a shape is endowed with a metric tensor at each point in the parametrization domain. Here, we consider a special canonical parametrization of the shape whose metric tensor at each point is Euclidean. This allows us to define a generic deformation descriptor, namely Canonical Distortion Coefficients (CDCs), which can be used to characterize surface deformations and can be efficiently computed in the discrete setting. Based on this, we formulate dense surface registration using a higher-order MRF with special topology, and develop an associated inference algorithm that requires minimal memory and achieves significant speedup via parallel acceleration. Last but not least, in order to reduce the search space, 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.

1.1 Related work

Modeling surface deformation is a very challenging task due to the large number of degrees of freedom that exist in real-world deformations. Certain “rigidity” assumptions have been widely made, 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.

Extrinsic surface representations were studied exhaustively for matching two surfaces undergoing a global rigid deformation (*i.e.*, rotation and translation) through the iterative closest points (ICP) [8] algorithm as well as its numerous variants (*e.g.*, [16]). Whereas global rigidity does not take into account bendable shapes (*e.g.*, garments or rubber bands), the notion of *local rigidity* has been proposed defining the deformation between two local corresponding neighborhoods as rigid [29]. However, searching for the correspondences between surfaces undergoing large deformations directly in the original extrinsic space may suffer from high computational complexity.

Intrinsic surface representations through distance functions and the expression of the surface matching problem as a volume registration one were studied initially to provide a dense solution to the correspondence problem (*e.g.*, [30, 45]). 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 (*e.g.*, [13, 14, 64, 68, 72]). Such representations usually require an initial set of feature correspondences or boundary conditions, which are difficult to find. The performance of these methods degrades in the presence of noise as well as varying scales, boundary conditions and resolutions. Furthermore, since most surface deformations are not perfectly isometric, solely considering intrinsic information introduces approximation errors to the matching result. In order to address the above-mentioned issues, [5, 54, 55] proposed to search for correspondences using a probabilistic formulation based on geodesic distances. Nevertheless, issues like the computational complexity and inaccuracy of geodesic distances towards establishing dense correspondences reduce the applicability of these methods. An alternative approach for seeking correspondences in the intrinsic space is to map the shape to a high dimensional space in which the Euclidean distance approximates the intrinsic properties of the surface. For example, the idea of diffusion maps [44, 50] is to represent the shape through a space that preserves the commute time, which has the advantage of being robust to topological changes. To overcome the high computational complexity issues in establishing point correspondences, the idea of functional maps [43] was proposed to establish the correspondences in the functional space, which is nevertheless limited to shapes with bijective mapping.

Graph matching is a powerful framework for establishing correspondences [3, 17], which is able to combine multiple matching costs (*e.g.*, appearance similarity and geometric compatibility) within an objective function through the integration of singleton, pairwise or higher-order interactions among nodes (*e.g.*, [24, 41, 58]). It has been employed often in the literature (*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 has been very limited, probably due to the fact that more sophisticated metrics are required in order to encode invariant spatial relationships between points on the surfaces [60]. Numerous optimization methods have been proposed in the context of graph matching, such as spectral relaxation (*e.g.*, [22, 40]), continuous relaxation (*e.g.*, [48, 57]), and randomized algorithms (*e.g.*, [39]). Notably, [58] proposed a novel pairwise graph-matching algorithm based on the dual-decomposition framework, which provides certain optimality guarantee on the solution.

Markov Random Fields (MRFs), in particular pairwise MRFs, have been widely applied to address numerous computational visual perception tasks, such as image segmentation, stereo, detection and registration, *etc.* [9, 61]. Recently, the applications of higher-order MRFs in computer vision have become more and more popular, mostly driven by the development of inference and learning methods (*e.g.*, [31, 36, 38, 66]). Higher-order MRFs, compared with pairwise models, allow a better characteri-

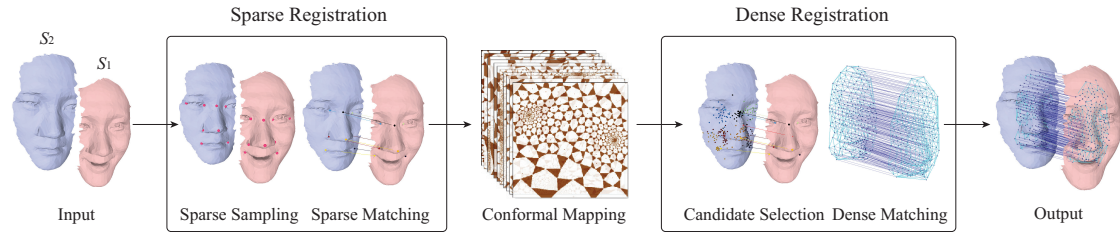


Figure 1: Overview of our algorithmic framework for surface registration as described in Sec. 1.2.

zation of statistics among random variables and have better expression power [61]. This is particularly useful for modeling measures that intrinsically involve more than two variables (*e.g.*, [4, 28, 69]) and characterizing invariant statistics (*e.g.*, [62, 63]).

1.2 Overview

A brief overview of the proposed surface registration system is shown in Fig. 1, where sparse surface matching is first performed and is followed by dense registration.

In the sparse registration stage, sparse feature points are selected according to geometric principles (such as the local maxima of Gaussian curvature [42] and the average geodesic distance function [33] between the input surfaces S_1 and S_2). On these points we find n_s correspondences between the two feature sets, via a higher-order graph matching algorithm that uses multiple matching criteria and is solved via the dual-decomposition technique (Sec. 2).

Once such correspondences have been established, the dense registration 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 provide $\binom{n_s}{3}$ candidate matching points on S_2 for any point $p \in S_1$. These points are then clustered to obtain meaningful matching candidates for each point (Sec. 3.2.2). 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). A higher-order MRF optimization is then formulated and solved to obtain the optimal dense registration result (Sec. 3.2, 3.3).

Finally, experimental validation (Sec. 4) and discussion (Sec. 5) conclude the paper.

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 potential correspondences between P_1 and P_2 . We introduce 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 $p_a^2 \in P_2$ are the two points defining a potential 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 missing correspondences are allowed in order to deal with partial matching. In such a graph labeling setting, the higher-order graph matching problem can be formulated as follows:

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

where θ_a is the matching cost for each correspondence $a \in P$, θ_{ab} for a pair of correspondences $(a, b) \in P \times P$, θ_{abc} for a triplet of correspondences $(a, b, c) \in P \times P \times P$, and $\boldsymbol{\theta}$ denotes the vector consisting of all such parameters. Note that in order for Eq. 2 to have non-trivial solutions, at least some values in $\boldsymbol{\theta}$ must be negative. Furthermore, the matching constraints can be reduced to pairwise terms in the energy function. More specifically, we observe the following equivalence:

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

where θ^∞ is a sufficiently large number. We use $P^C = \{(a, b) \mid a, b \in P, a \neq b \text{ and } (p_a^1 = p_b^1 \text{ or } p_a^2 = p_b^2)\}$ 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 a pseudo-boolean optimization problem [10] as follows:

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

The above formulation is generic 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 its optimization is an NP-hard problem [10]. We adopt the flexible dual-decomposition technique [36] to perform the optimization efficiently. Moreover, theoretically any higher-order terms can be efficiently reduced into quadratic terms [31], often referred to as *order reduction*. The basic idea is to transform an energy function consisting of higher-order terms into an equivalent one that has the same minimum but involves only singleton and pairwise terms.

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 encode both appearance and geometric similarities, the pairwise terms constrain the matching solution space and the higher-order terms encode intrinsic deformation errors.

2.2.1 Singleton potentials

For each correspondence, we consider both geometric and texture information to define its potential [56]. Let us denote the Gaussian curvature at point p as $\text{curv}(p)$, and the texture value at point p as $\text{tex}(p)$. The singleton potential for a correspondence $a = (p_a^1, p_a^2)$ 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.

2.2.2 Pairwise potentials

In the general higher-order graph matching formulation as shown in Eq. 4, we use pairwise potentials to encode the mapping constraints for the graph matching, by setting θ^∞ and θ_{ab} to be 10^5 and 0, respectively, in our experiments. Other pairwise potentials for surface matching, such as those of [60], can also be incorporated straightforwardly in this formulation.

2.2.3 Higher-order potentials

The uniformization theorem [25] states any 3D surface can be flattened conformally to a canonical 2D domain. Such a mapping represents a feature point p as a parametric coordinate in the complex plane $z_p \in \hat{\mathbb{C}}$. Conformal mappings are flexible because of the 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 [42], we compute the matching score between two triplets as the deformation error based on the Möbius transform.

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 maps 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 endow 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 [42], 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:

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

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 deal with non-plausible correspondences (in our experiments $\delta = 0.1$). Intuitively, the more matching pairs and the smaller the distance between those mutually closest pairs, the lower the potential energy.

However, the Möbius energy introduces ambiguity since it assumes isometry, which is invariant under symmetric transformations (an example is shown in Fig. 2). To resolve such an ambiguity, we consider an

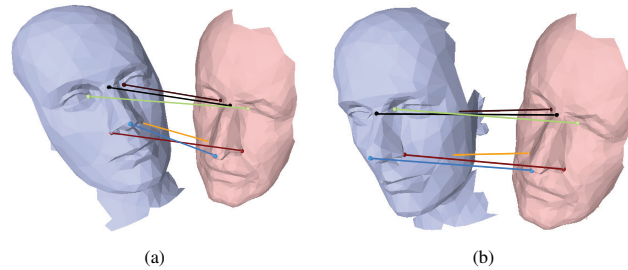


Figure 2: An example showing the matching ambiguity when considering only intrinsic information. The matching scores in (a) and (b) are the same according to Eq. 6 based on the Möbius transform, since the distances between the matched features are identical. However, such ambiguity can be avoided by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

extrinsic property of the surface, *the Gaussian map*, defined as the mapping of the normal at each point on the surface to the unit sphere [27]. Two triplets have the same orientation if and only if the determinants of their normals are of the same sign. Therefore, in order to avoid ambiguities in orientation, a constraint on the Gaussian map is imposed via the following potential:

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

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 considered as a soft constraint in our framework, since normals can change orientations when the surface undergoes very large deformations.

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 defining the potentials of the graph matching problem (Eq. 4), 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 [35, 36]. For the graph matching problem, let $\boldsymbol{\theta}$ denote the vector of the weights of the singleton, pairwise and triplet terms, and I denote the set of subproblems. The decomposition is expressed as $E(\mathbf{x}|\boldsymbol{\theta}) = \sum_{\sigma \in I} \rho_{\sigma} E^{\sigma}(x|\boldsymbol{\theta}^{\sigma})$ where ρ_{σ} is the weight for each subproblem. Then the original problem is solved by updating the parameter $\boldsymbol{\theta}^{\sigma}$ of each subproblem σ so that it increases the energy of the dual problem, while satisfying the following decomposition constraint:

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

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

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

This lower bound is maximized using a projected subgradient method so that a solution to the original problem can be extracted from the Lagrangian solutions [36].

Specifically, we decompose the optimization problem in Eq. 2 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 using the Hungarian algorithm [3].
2. a **higher-order pseudo-boolean subproblem** by reducing the higher-order terms to quadratic terms [10] which can be solved by the QPBO algorithm [34]. Here we employ [31] for the reduction. Specifically, each third-order term $\theta_{abc} x_a x_b x_c$ in Eq. 4 is replaced by $\theta_{abc} w(x_a + x_b + x_c - 2)$ if $\theta_{abc} < 0$, and $\theta_{abc} \{w(x_a + x_b + x_c - 1) + (x_a x_b + x_b x_c + x_c x_a) - (x_a + x_b + x_c) + 1\}$ otherwise (where w is an auxiliary binary variable), leading to an objective function with the same minimum as the original one.
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.

We introduce a higher-order pseudo-boolean subproblem, in contrast to [58] that considers pairwise subproblems. Given the solutions of the subproblems, we update the dual variables $\{\theta^\sigma\}$ by projecting them to the space that satisfies Eq. 9 [36, 58]. This process is performed iteratively until convergence.

The above algorithm involves an expensive step in each iteration, which is the max-flow computation. We employ the implementation of [37], 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.

3 Dense surface registration using higher-order MRFs

The main obstacles in extending the sparse matching approach to dense relate to:

- *Model accuracy*: the Möbius matching potential defined in Eq. 7 has high discriminative power when the feature points are far apart. However, such power degrades as the sampling becomes denser, which affects the optimality properties of the obtained solution. Hence, we require an alternative potential that is able to encode localized and anisometric deformations.
- *Computational complexity*: increasing the number of sampling points n would make the sparse matching approach computationally prohibitive, since the graph structure would grow in the order of $O(n^6)$ if we consider all possible triplets.

These issues are addressed through a novel higher-order MRF-based dense surface registration method that is endowed with: (i) a new deformation model; (ii) a novel graph-based energy to determine the optimal instance of this deformation model; (iii) an efficient matching candidate selection scheme; (iv) an efficient optimization algorithm for the MRF inference.

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 living on a surface. We first introduce the generic deformation model in the continuous setting and then derive its variant for the discrete setting.

3.1.1 Continuous setting

Let $(\mathcal{M}, g_{\mathcal{M}})$ denote a surface \mathcal{M} endowed 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:

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

Different local representations describe the same surface, if the following chain rule is 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 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 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*. We will show that considering the canonical parametrization/Jacobian allows us to characterize arbitrary deformations between surfaces independently of the choice of intrinsic/extrinsic surface representations, which is the main advantage of our deformation model.

Let us 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 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 .

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. 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 $f : U_\alpha \rightarrow U_\beta$, between any canonical parametrization x_α and x_β for p and q respectively, we define

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

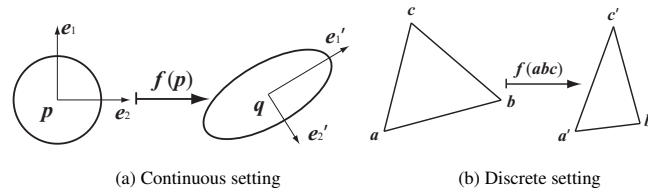


Figure 3: The finite element method assumes the transformation between facets to be piecewise linear and $f(\vec{ab}) = a'\vec{b}'$, $f(\vec{ac}) = a'\vec{c}'$. Under the linearity assumption, the Jacobian can be computed in a closed form for each pair of triangular facets $\triangle abc \mapsto \triangle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

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 provides 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 useful 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 hyper dimensions (nD).

3.1.2 Discrete setting through finite elements

In the context of finite element analysis [12], a continuous space is 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 popular 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.

In the continuous setting, the Jacobian matrix between p and q is a linear transformation that maps the tangent spaces at p and q . 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, $J(\vec{bc}) = b'\vec{c}'$ is satisfied, *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. For n -manifold shapes, the computation of CDCs only requires solving n linear equations and eigenvalues. Note that the computation looks analogous to the surface parametrization of [46, 47], due to the piecewise linear assumption. However, Alg. 1 is derived in the context of Riemannian geometry for shape deformation.

Based on CDCs, we can deform the original shape (*e.g.*, [52]), and also determine the correspondences between two shapes. We now introduce a general MRF formulation for dense shape registration.

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 $\vec{a'b'}$ and \vec{ac} to $\vec{a'c'}$.

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

Step Four: Output λ_1 and λ_2

3.2 MRF formulation for shape registration

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 , shape registration seeks to determine the correspondence for each point $p \in \mathcal{V}$ on S_2 , which we obtain via inference over a higher-order MRF framework.

The considered MRF model is a hyper-graph $\mathcal{G} = (\mathcal{V}, \mathcal{F})$, corresponding to a triangulation of 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 .

The singleton potential $\theta_u(x_u)$ of the objective function is defined as the difference in the feature descriptor (e.g., texture, 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 $\triangle uvw$ to $\triangle x_u x_v x_w$ (Alg. 1), the higher-order potential can then be defined as:

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

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 refers to the 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 addressing the optimization of the objective function.

3.2.1 Deformation constraints

The assumption that *deformations are similar across different shapes of the same type* is valid in the context of natural shapes, and has been applied in, for example, deformation transfer [53] and facial expressions transfer [59]. 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 3D data, shown in (a) and (c), are captured with markers using the system described in [65]. 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:

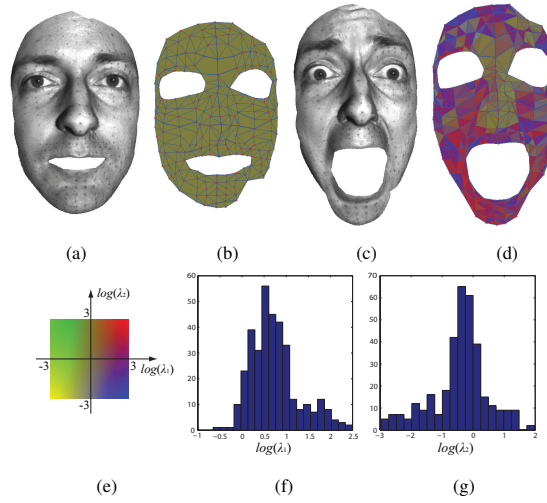


Figure 4: An expression deformation prior obtained by 3D scanned data with markers. (a) and (c) show the 3D scan of the onset and peak of a facial expression with large shape deformations respectively. (b) and (d) are the corresponding triangular templates constructed from the 3D scan data. (d) shows 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).

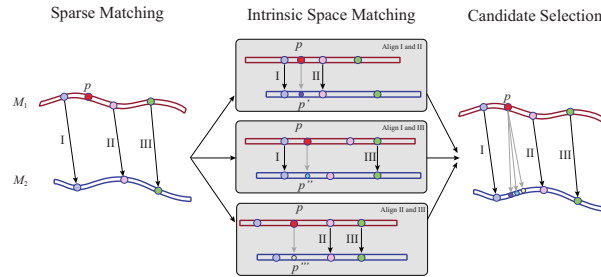


Figure 5: 2D Illustration of candidate selection scheme.

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

where λ_1 and λ_2 are the CDCs obtained by matching $\triangle uvw$ to $\triangle x_u x_v x_w$.

3.2.2 Matching candidate set

Inclusion of all the points on S_2 as the candidate set of 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 large deformation. Therefore, the neighborhood relations of 3D points cannot be straightforwardly defined based on their 3D coordinates. However, given *three correspondences*, an alignment of the two surfaces in the conformal mapping domain can be obtained in a closed form by determining the associated Möbius transformation. The alignment is accurate when the two surfaces undergo perfect isometric deformation and the error increases continuously as the deformation deviates from isometry. Hence, the

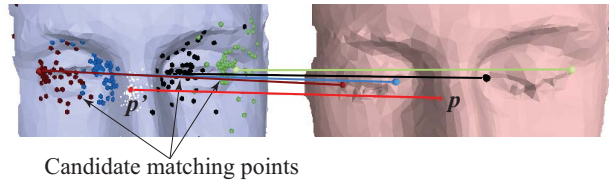


Figure 6: An example showing candidate points obtained from different Möbius transforms and their clustering. For any point p from the source surface, the clustering of its matching candidate points on the target surface gives us a matching candidate p' .

matching candidate set for dense surface registration can be obtained from the sparse registration results (Fig. 1), which we achieve via two successive operations: candidate selection and clustering.

Given 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 . Triplets of sparse correspondences are mapped to a prefixed configuration by solving a Möbius transformation, as described in Sec. 2.2. A pair of corresponding triplets in the 2D domain produce a matching of the two surfaces and provide a corresponding 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, where we observe that most candidate points are close to the desired correspondences. It should be noted that considering all triplet correspondences provides an exhaustive set of matching candidates. We further reduce the search space by clustering such candidates using mean shift [21] and maintain only the principal modes of the density (other approaches such as [19] and [20] can be used as well).

3.3 An efficient higher-order MRF optimization

Inference of higher-order MRFs is an active research topic and various techniques have been proposed. Most existing approaches either employ *order-reduction* [31, 26] (first reduce higher-order terms into pairwise ones and then solve the problem using *graph cuts* techniques [34]) or *dual-decomposition* [35]. However, the algorithms designed for general MRFs often lack efficiency in terms of computation and/or memory, and fail to take advantage of MRFs with special topology and/or potentials. In order to efficiently perform inference, we explore the topological property of our MRF and develop a parallel optimization algorithm that requires minimal memory and achieves significant speedup via GPU.

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:

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

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:

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

Here $M_{uvw;u;i}$ is the dual variable (message) corresponding to the constraint $\sum_{j,k} \tau_{uvw;ijk} = \tau_{u;i}$ (Fig. 8(a)).

The dual problem of Eq. 17 can be solved by *min-sum diffusion* algorithm [66] as shown in Alg. 2. It has been shown that once convergence is attained, the solution satisfies the J -consistency condition [66]. Furthermore, each update of the message requires a simple reparameterization of the MRF, and does not need extra memory for storing all the dual variables $M_{uvw;u;i}$. Hence, the memory requirement for Alg. 3 is only for storing primal variables, *i.e.*, $O(|\mathcal{V}||\mathcal{L}| + |\mathcal{F}||\mathcal{L}|^3)$.

Algorithm 2: Min-sum diffusion algorithm.

```

repeat
  for each  $M_{uvw;u;i}$  do
     $M_{uvw;u;i} \leftarrow \frac{1}{2}[\theta_{u;i} - \min_{j,k} \theta_{uvw;ijk}]$  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 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.

To explore the parallelism of the min-sum algorithm (Alg. 2), we define the concept of independent facet set.

Definition 3. (Independent facet set) Given a graph $\mathcal{G} = (\mathcal{V}, \mathcal{F})$, a subset $\mathcal{F}_k \subset \mathcal{F}$ is called an independent facet 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 facet 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 facet sets $\cup_i \mathcal{F}_i$

repeat

for each Independent facet 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 reparameterize (Alg. 2).

end for

until convergence

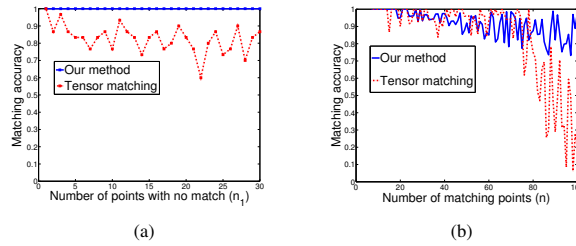


Figure 7: Comparisons between our algorithm and the tensor matching algorithm [24]. (a) shows the performance in partial matching when $n = 30$ points is matched to $n + n_1$ points. Our method is strongly robust to outliers. (b) shows the use of partial pairwise terms for matching constraints to overcome memory limit for full matching ($n_1 = 0$) and our method remains robust with increasing n .

4 Experimental results

Our system is tested on an Intel[®] Xeon(TM) 3.4G PC with 4G RAM and an NVIDIA[®] Geforce 9800GTX+ graphics card. We first present the experimental evaluation on our higher-order graph matching and higher-order MRF optimization algorithms, and then exhibit the performance of the whole method on sparse & dense shape registration, and surface tracking.

4.1 Higher-order graph matching

To evaluate the performance of our higher-order graph matching algorithm, in particular its ability to handle partial matching, we compare our algorithm with the tensor-based algorithm of [24] using the authors' implementation. In order to easily compare the algorithms quantitatively, we follow the synthetic data generation of [24]. Given $n > 0$, we randomly sample n points on the 2D plane to obtain the first point set S_1 , and then generate their matching points by randomly rotating, translating and scaling those points in S_1 plus a random noise, resulting in a second point set S_2 . In addition, in order to evaluate the performance of partial matching, we enlarge S_2 by including $n_1 > 0$ randomly generated points, which have no correspondence (*i.e.*, $|S_1| = n$ and $|S_2| = n + n_1$). Similar to [24], we set singleton potentials as zero, and use the following similarity-invariant higher-order potentials:

$$\theta_{abc} = A e^{-((\angle p_a^1 - \angle p_a^2)^2 + (\angle p_b^1 - \angle p_b^2)^2 + (\angle p_c^1 - \angle p_c^2)^2)},$$

where $\angle p_a^k$, $\angle p_b^k$ and $\angle p_c^k$ denotes the vertex angles of $\triangle p_a^k p_b^k p_c^k$, and the weight constant A is set to be -100 in our algorithm. We sample $50n$ triangles for the triplet potentials and use the same ANN-based algorithm for computing the tensor as in [24]. Fig. 7(a) shows the matching accuracy with increasing n_1 ($n = 30$). Note that our algorithm remains robust even when $n_1 = n$, which demonstrates its ability to handle partial matching.

One limitation of our algorithm is the large memory requirement for encoding those pairwise potentials that ensure a valid matching result (Eq. 3). If we consider all the constraints in Eq. 3, the algorithm is

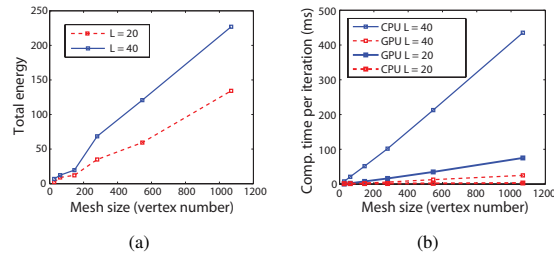


Figure 8: Performance analysis of our MRF optimization algorithm. (a) shows the optimality using the test cases described in Sec. 4.2. (b) shows the speedup using the parallel implementation of Alg. 3. L is the number of labels for each node. We show the runtime per iteration since different inputs have different iteration counts.

only able to handle 30 – 40 matching pairs. However, this can be partially overcome by only considering a fraction of the pairwise constraints and selecting matching results that are valid (*i.e.*, Eq. 1 is satisfied). Fig. 7(b) shows the percentage of correct matching as n increases ($n_1 = 0$), by limiting the maximal number of pairwise terms to be $100n$. We observe that our algorithm degrades less sharply than [24] even in the case of full matching ($n_1 = 0$). However, this gain in accuracy does not come for free, our algorithm runs 2 – 10 times slower than that of [24] on average.

4.2 Higher-order MRF optimization

We implement Alg. 3 using the NVIDIA[®] CUDA architecture [32]. 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

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

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

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

In this case, the optimal solution of Eq. 15 should be $(l_u)_{u \in \mathcal{V}}$. Fig. 8(a) shows the result of our algorithm using the above designed test cases for different mesh and label sizes. 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. 8(b) shows the comparison on average time taken per iteration, between the implementations with and without GPU accelerations. The number of iterations depends on the (unknown) form of the objective function. In our experiments, the algorithm converges within 3000 iterations.

4.3 Sparse & dense shape registration

In this section, we evaluate our approach for dense surface matching. Let us first introduce some related implementation details and then discuss the obtained results.

In sparse surface matching, for the potential functions of the surface 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 $\lambda_2 = 1$. For each input mesh, one conformal mapping is computed using the mid-edge uniformization algorithm [42, 46], which involves

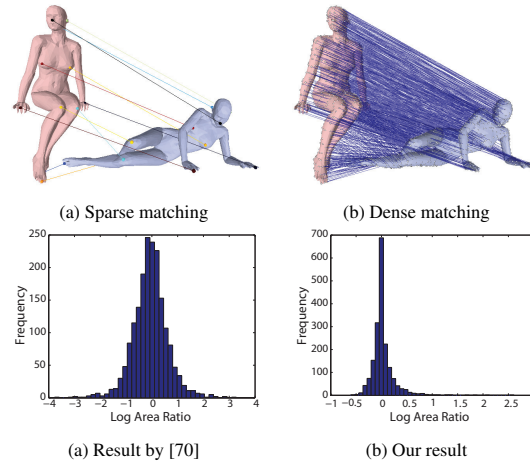


Figure 9: Matching result for the body data: (matched/total = 2861/3376)

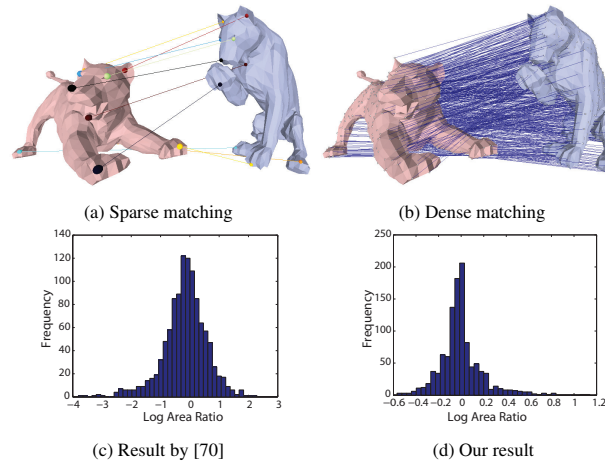


Figure 10: Matching result for the lion data: (matched/total = 1105/1251)

solving a symmetric linear equation (it takes $< 1s$ for a mesh with 10^4 faces using a GPU implementation). Additional conformal mappings are computed in a closed form by the Möbius transformation induced by three correspondences (Sec. 2.2).

The candidate selection and clustering of 10^3 points based on 10 sparse features takes approximately 1 minute (as described in Sec. 3.2.2). The input to the dense surface registration stage is the set of vertices \mathcal{V} on the source surface with at least one valid matching point from our candidate selection stage. 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 $v \in \mathcal{V}$. We then uniformly re-sample $L = 64$ points for each $v \in \mathcal{V}$ near the original matching candidates, to further improve matching accuracy. The computation of all the L^3 possible CDCs for one facet takes only 2.0ms on average using the GPU. The computation of the potential $\theta_{uvw;ijk}$ for a graph with $|\mathcal{F}| = 2000$ takes 3s.

We consider the challenging problem of matching surfaces that involve large deformations and inconsistent boundaries (partial overlapping). The number of vertices for each mesh is in the range of

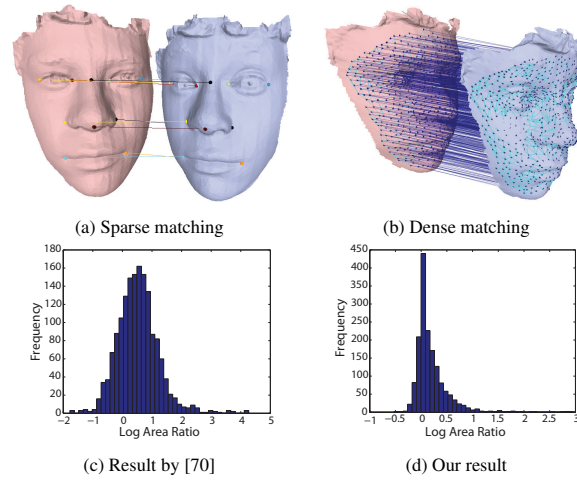


Figure 11: Matching result for the face data: (matched/total = 2098/2644)

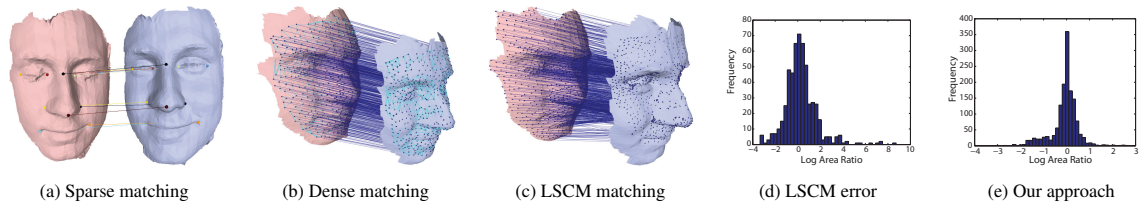


Figure 12: Comparison with LSCM approach [64] for dense surface matching. (matched/total = 1455/1635) (best viewed in color).

1,500 – 4,000. Our method produces 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). The lion data of Fig. 10 comes from [53] and the face and hand data are captured with texture by the 3D scanner described in [65]. Based on 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, to measure the quality of the dense registration. The local area is not expected to undergo abrupt change in natural deformations (*e.g.*, expression change, stretched arms or bending figures). 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 is able to map the boundaries of the surface to slits while preserving the conformal structure of the surface in an exact sense. Therefore, 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 Figs. 11, 12, and 13. An example of significant non-overlap between the two meshes is shown in Fig. 1.

Matching with large deformations: Figs. 9, 10, 13 and 14 show results matching two surfaces undergoing a large deformation. Even when the sparse features cannot all be selected consistently (as shown in Fig. 14), our higher-order graph matching algorithm in Sec. 2.3 is able to find reliable sparse correspondences (Fig. 14(a)) and obtain a dense surface matching result through the two-stage optimization scheme in Sec. 1.2 (Fig. 14(b)).

Comparison experiments: Fig. 12 shows a comparison between our algorithm and the Least Square Conformal Mapping (LSCM) approach [64]. Although LSCM can handle free boundaries, there is no

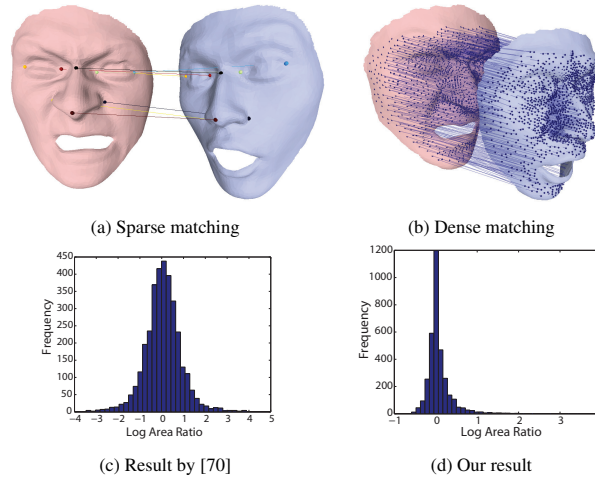


Figure 13: Dense matching under large non-rigid deformations. (matched/total = 2378/3633)

theoretical guarantee that the conformal structure is preserved near the boundary and it can produce self-intersections in the mapping space [51]. In our comparison, we use the feature correspondences computed from the sparse matching stage as the feature constraints for LSCM. The limitations of LSCM can be observed in Fig. 12(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 we cannot compare directly with the results in [64], since their initial feature points were manually selected.

Furthermore, we compare our results with a more recent dense surface matching approach [33], using the authors’ implementation. Similar to [33], 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 f denotes the correspondence map from S_1 to S_2 obtained by an matching algorithm, f_{true} the ground-truth correspondence map, and $d_{S_2}(\cdot, \cdot)$ 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 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 [33] is shown in Table 1. Note that [33] 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.

4.3.1 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 [71]. 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 [71] 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 $I_2 = [0.846, 1.182]$ for λ_1 and λ_2 respectively (Eq. 16). Furthermore, we handle drift errors by imposing

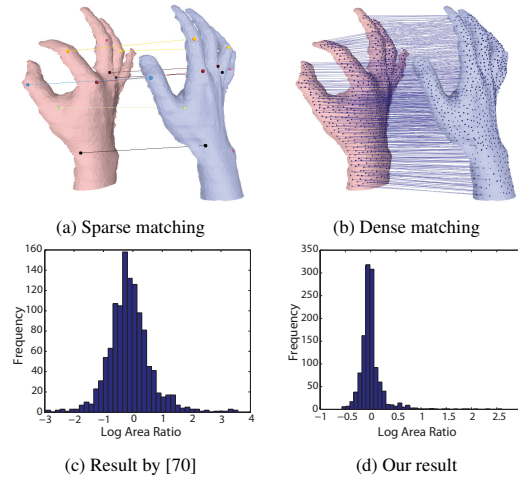


Figure 14: Dense matching under multiple articulated deformations. (matched/total = 1224/1786)

Data	Our method	Kim <i>et.al.</i> [33]
Body (Fig. 9)	0.0622	0.2437
Lion (Fig. 10)	0.0832	0.1790
Face (Fig. 11)	0.0319	0.0465
Face (Fig. 13)	0.0565	0.0865
Hand (Fig. 14)	0.0481	0.1193

Table 1: Comparison with a recent intrinsic method for dense surface registration [33]. The average error is calculated based on Eq. 18.

consistency between the first frame and the current frame, using the same deformation prior obtained in Fig. 4.

Fig. 15 shows the tracking results on the BU-4DFE database [67], which consists of 3D dynamic facial expressions of different subjects. A mesh template is manually constructed in the first frame and automatically tracked in the subsequent frames. Because of the temporal continuity in consecutive frames, sufficient matching candidates can be obtained by only looking at the 3D 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 accurate tracking results for sequences with significant anisometric facial deformation as shown in Fig. 15.

5 Conclusion

We proposed a higher-order graph-based approach for dense, non-rigid surface registration. Specifically, a two-stage algorithm was introduced to reduce 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 combined the similarities in appearance and geometry as well as the intrinsic deformation error based on the isometry assumption. In our dense surface registration stage, a generic deformation model, introduced to a higher-order MRF formulation, was proposed to handle anisometric surface deformations. The proposed method achieved

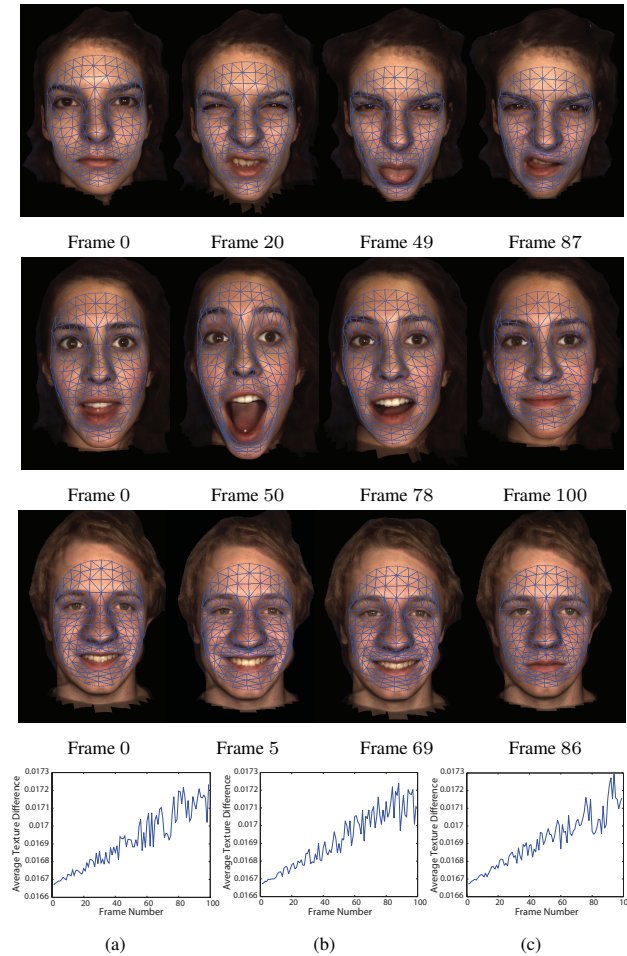


Figure 15: Shape tracking results. The average texture difference for all correspondences between every frame and the first frame for the three sequences, from top to bottom are shown in (a), (b) and (c), respectively.

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 our higher-order MRF optimization. We will also explore more sophisticated ways to learn the CDCs and apply them in expression transfer, face recognition, *etc.*

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.
- [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] R. Burkard, M. Dell’Amico, and S. Martello. *Assignment Problems*. Society for Industrial and Applied Mathematics, 2009.
- [18] 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.

- [19] M. Cho, J. Lee, and K. M. Lee. Feature correspondence and deformable object matching via agglomerative correspondence clustering. In *Proc. IEEE Int'l Conf. Computer Vision*, 2009.
- [20] M. Cho and K. M. Lee. Progressive graph matching: Making a move of graphs via probabilistic voting. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2012.
- [21] 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.
- [22] T. Cour, P. Srinivasan, and J. Shi. Balanced graph matching. In *Proc. Neural Information Processing Systems*, 2007.
- [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] 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.
- [30] X. Huang, N. Paragios, and D. N. Metaxas. Establishing local correspondences towards compact representations of anatomical structures. In *International Conference, Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2003.
- [31] H. Ishikawa. Higher-order clique reduction in binary graph cut. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2009.
- [32] A. Keane. CUDA (compute unified device architecture). 2006.
- [33] V. G. Kim, Y. Lipman, and T. Funkhouser. Blended intrinsic maps. *ACM Trans. Graph. (Proc. of SIGGRAPH 2011)*, 30:79:1–79:12, 2011.
- [34] 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.
- [35] N. Komodakis and N. Paragios. Beyond pairwise energies: Efficient optimization for higher-order MRFs. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2009.
- [36] 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.
- [37] 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.

- [38] N. Komodakis, B. Xiang, and N. Paragios. A framework for efficient structured max-margin learning of high-order MRF models. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 37(7):1425–1441, 2015.
- [39] J. Lee, M. Cho, and K. M. Lee. Hyper-graph matching via reweighted random walks. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2011.
- [40] 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.
- [41] M. Leordeanu, A. Zanfır, and C. Sminchisescu. Semi-supervised learning and optimization for hypergraph matching. In *Proc. IEEE Int'l Conf. Computer Vision*, 2011.
- [42] Y. Lipman and T. Funkhouser. Möbius voting for surface correspondence. *ACM Trans. Graph.*, 28(3):72:1–72:12, 2009.
- [43] 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.
- [44] 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.
- [45] N. Paragios, M. Rousson, and V. Ramesh. Non-rigid registration using distance functions. *Computer Vision and Image Understanding*, 89(203):142–165, 2003.
- [46] U. Pinkall and K. Polthier. Computing discrete minimal surfaces and their conjugates. *Experimental Mathematics*, 2(1):15–36, 1993.
- [47] P. V. Sander, J. Snyder, S. J. Gortler, and H. Hoppe. Texture mapping progressive meshes. In *SIGGRAPH '01*, pages 409–416, 2001.
- [48] C. Schellewald and C. Schnorr. Probabilistic subgraph matching based on convex relaxation. In *Energy Minimization Methods in Computer Vision and Pattern Recognition (EMMCVPR)*, 2005.
- [49] 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.
- [50] 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.
- [51] A. Sheffer, E. Praun, and K. Rose. Mesh parameterization methods and their applications. *Found. Trends. Comput. Graph. Vis.*, 2(2):105–171, 2006.
- [52] O. Sorkine and M. Alexa. As-rigid-as-possible surface modeling. In *Symposium on Geometry Processing*, pages 109–116, 2007.
- [53] R. W. Sumner and J. Popović. Deformation transfer for triangle meshes. In *SIGGRAPH*, pages 399–405, 2004.
- [54] 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.
- [55] 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.

- [56] N. Thorstensen and R. Keriven. Non-rigid shape matching using geometry and photometry. In *Proc. Asian Conf. Computer Vision*, 2009.
- [57] P. H. S. Torr. Solving Markov Random Fields using semi definite programming. In *Ninth Int'l Workshop on Artificial Intelligence and Statistics*, 2003.
- [58] 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.
- [59] D. Vlasic, M. Brand, H. Pfister, and J. Popović. Face transfer with multilinear models. *ACM Trans. Graph.*, 24:426–433, 2005.
- [60] 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.
- [61] 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.
- [62] 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.
- [63] 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.
- [64] 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.
- [65] 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.
- [66] 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.
- [67] 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.
- [68] 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.
- [69] 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.
- [70] 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.

- [71] 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.
- [72] D. Zhang and M. Hebert. Harmonic maps and their applications in surface matching. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 1999.

Contents

1	Introduction	3
1.1	Related work	4
1.2	Overview	5
2	Sparse surface registration using higher-order graph matching	5
2.1	Pseudo-boolean higher-order graph matching	5
2.2	Higher-order graph matching for sparse surface matching	6
2.2.1	Singleton potentials	7
2.2.2	Pairwise potentials	7
2.2.3	Higher-order potentials	7
2.3	Optimization and computational complexity	8
3	Dense surface registration using higher-order MRFs	9
3.1	A generic deformation model	9
3.1.1	Continuous setting	9
3.1.2	Discrete setting through finite elements	11
3.2	MRF formulation for shape registration	12
3.2.1	Deformation constraints	12
3.2.2	Matching candidate set	13
3.3	An efficient higher-order MRF optimization	14
4	Experimental results	16
4.1	Higher-order graph matching	16
4.2	Higher-order MRF optimization	17
4.3	Sparse & dense shape registration	17
4.3.1	Surface tracking	20
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