

As-Conformal-As-Possible Surface Registration

Yusuke Yoshiyasu¹, Wan-Chun Ma², Eiichi Yoshida^{1,3} and Fumio Kanehiro¹

¹National Institute of Advanced Industrial Science and Technology (AIST), Japan

²Victoria University of Wellington, New Zealand

³CNRS-AIST JRL (Joint Robotics Laboratory), UMI3218/CRT

Abstract

We present a non-rigid surface registration technique that can align surfaces with sizes and shapes that are different from each other, while avoiding mesh distortions during deformation. The registration is constrained locally as-conformal-as-possible such that the angles of triangle meshes are preserved, yet local scales are allowed to change. Based on our conformal registration technique, we devise an automatic registration and interactive registration technique, which can reduce user interventions during template fitting. We demonstrate the versatility of our technique on a wide range of surfaces.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Computer Graphics—Computational Geometry and Object Modeling

1. Introduction

Recently, there has been growing interest in creating high-quality digital three-dimensional (3D) models from the 3D scans of real-world objects, driven by demand from industries such as films and manufacturing, e.g., digital doubles of actors for visual effects and digital 3D models for personal fabrication using 3D printing technology. Progress in 3D scanning allows us to capture high-resolution and highly-detailed 3D geometries even for dynamic scenes. Unfortunately, these 3D scans are not readily usable in real-world applications because scanned surfaces are partial, incomplete and contain holes/noises. In many applications, a complete mesh model with a good topology (neatly aligned connectivity edges) is desirable because it is easier to manipulate and attach texture and bump maps.

Surface registration is thus an inevitable post-process of 3D scanning, which aligns and merges captured surfaces in order to generate a complete surface model. An effective way to do this is the technique called template-based model fitting, which starts from a template mesh and gradually deforms it toward the scan. By fitting a template with good topology, we can obtain a high quality mesh model that can be used in a wide range of digital applications. Nonrigid surface registration is also vital when adopting mesh-based re-targeting [SP04], which demands correspondences between the source model and the target, and blendshape techniques




Figure 1: As-conformal-as-possible (ACAP) nonrigid registration. Given only five correspondences (forehead, cheeks and ears), our technique is able to fit the template to the highly-detailed scan with different size. Notice that the result shows high-quality mesh connectivities with no fold-overs and less shear distortions.

[WLP09], which require blendshape models with identical connectivities.

While many techniques have been proposed, non-rigid registration is still a problem, particularly because of the following challenges:

Geometric and semantic consistency Because the orientations, sizes, shapes and poses of two surfaces are often not equivalent, aligning the models globally while also capturing surface details at fine scale is difficult. Furthermore, semantic consistency is crucial for nonrigid registration. For example, in the case of face registration, features around the eyes, lips and ears must correspond to each other.

Mesh-connectivity preservation During the registration of two surfaces with drastically different geometries, the mesh undergoes large deformation. In such a situation, the mesh is susceptible to shear distortions and fold-overs. Obtaining a clean and high-quality mesh that is usable for applications is very challenging.

Less user effort The registration technique should not require a large amount of user effort in order to specify many feature points by hand. However, due to the above mentioned two challenges, this is difficult to achieve; previous studies require specifying dozens of landmarks [SP04, YLSL10].

Early nonrigid registration techniques use smoothness regularization [ACP03, SP04, ARV07] to constrain deformations. This approach is flexible enough to align semantically different shapes such as sphere to tooth [YLSL10]. Nevertheless, it loses most of the original details and is also poor at preserving mesh structures during large deformation, which require specifying many correspondences by hand to work correctly. In contrast, regularization based on isometric (length-preserving) mapping [LSP08, HAWG08] provides strong constraints on deformations, which is often employed in automatic registration. However, this is not suitable for aligning surfaces with different sizes or shapes. Most often, the template mesh and scans share semantics, but are different in global scale, pose, local scale, and details; in other words, two surfaces are in the same class, e.g., different human individuals.

In this paper, we present the as-conformal-as possible (ACAP) surface registration technique, which is flexible enough to non-rigidly align two surfaces with different details and sizes, yet is able to preserve the mesh structure of the template. The key idea of our work is to incorporate conformal mapping into a surface registration framework in order to preserve the angles of triangles in the mesh. Unlike the techniques based on isometric mapping, this approach allows us to adjust the local scale of the model freely and at the same time maintain mesh structures. Furthermore, unlike the techniques based on smoothness regularization, this approach can produce a good initial shape from a small amount of initial correspondences for user intervention to be minimal. Conformal mappings have been used in parameterization [LPRM02] and correspondence algorithms [LF09], but they are often used for flattening the surfaces, which restricts the surfaces to be near-zero genus (or at least the surfaces must have the same genus). We, in turn, define a conformal energy directly in 3D space, enabling registration of incomplete surfaces with complex topology.

The rest of the paper is organized as follows. We first review related studies briefly in Section 2. Section 3 describes our core algorithm—ACAP surface registration. Next, we introduce a coarse-to-fine fitting strategy for efficient registration in Section 4. In Section 5, based on the ACAP framework, we propose an automatic registration technique beyond isometric deformations and an interactive registration

technique that requires only a small set of correspondences. The ACAP surface registration technique is tested on various types of scan data, including different expressions and poses (Section 6). We also compared ACAP with other state-of-the-art techniques.

In summary, our main contributions are the following:

- We derive a nonlinear conformal stiffness term and incorporate it into the nonrigid registration framework such that we can handle spatially varying scale changes of the template and the scan, while preserving mesh structures during deformation. We also introduce a linear formulation of the conformal stiffness term. Although it produces distortions when the mesh undergoes large deformations, the linearized version can be minimized more efficiently. Using the linear conformal stiffness, we improve the performance of our technique during the last stage of registration, where only small deformations are required.
- We propose an automatic registration technique employing a nonrigid shape matching technique that is robust to non-isometric deformations and data incompleteness. To the best of our knowledge, this is the first automatic nonrigid registration technique that can handle large non-isometric deformations.
- We propose an interactive registration technique that requires only a small amount of user interaction. Because conformal stiffness can preserve the overall geometry of the template while allowing scale, we can provide a good initial shape with at least three correspondences. The user can then specify the correspondences on-the-fly by observing the deformed model until a satisfactory result can be obtained. Thus, user interventions can be minimal.

2. Background

2.1. Classes of mappings

Here, we briefly review the classes of mappings [FH05] as a guide to choose an appropriate mapping for nonrigid surface registration. Suppose that S is a surface and f is a mapping from S to a second surface \tilde{S} . We consider a 3D-to-3D mapping case, where a point at position \mathbf{S} is deformed to $\tilde{\mathbf{S}}$ by a nonlinear function f , $\tilde{\mathbf{S}} = f(\mathbf{S})$. We define an orthogonal local frame $d\mathbf{S} = [d\mathbf{S}_1, d\mathbf{S}_2, d\mathbf{S}_3]$ at \mathbf{S} , which is deformed to $d\tilde{\mathbf{S}} = [d\tilde{\mathbf{S}}_1, d\tilde{\mathbf{S}}_2, d\tilde{\mathbf{S}}_3]$. A 3×3 local linear transformation \mathbf{T} called the deformation gradient is calculated as $\mathbf{T} = d\tilde{\mathbf{S}} \cdot (d\mathbf{S})^{-1}$. The rotation-invariant measure of deformation, the Cauchy-Green stretch tensor, is defined as $\mathbf{C} = \mathbf{T}^T \mathbf{T}$, which is an analogue of the first fundamental form. The properties of mappings are described as follows (Table 1).

Isometric mappings A mapping from S to \tilde{S} is isometric or length-preserving if the length of any arc on \tilde{S} is the same as that of the corresponding arc on S . When a mapping is isometric, \mathbf{C} is an identity matrix, $\mathbf{C} = \mathbf{I}_d$. In other words, deformation is locally rigid (rotation), $\mathbf{T} = \mathbf{R}$.

Conformal mappings A mapping from S to \tilde{S} is conformal

mal or angle-preserving if the angle of intersection of every pair of the intersecting arcs on \tilde{S} is the same as that of the corresponding arcs on S . When a mapping is conformal, the axes of the local frame must be orthogonal and have the same norm. In terms of stretch, it must satisfy $\mathbf{C} = s^2 \mathbf{I}_d$, where s is a scale. In other words, a local transformation is scale and rotation $\mathbf{T} = s\mathbf{R}$; i.e., it is a similarity transformation. Thus, the circle and the sphere are transformed to the circle and the sphere, but their radii are allowed to change from their original values.

Equiareal mappings A mapping from S to \tilde{S} is equiareal if every part of S is mapped onto a part of \tilde{S} with the same area. It is scale preserving.

Harmonic mappings A mapping from S to \tilde{S} is harmonic if the deformation minimizes the Dirichlet energy:

$$E_D(f) = \frac{1}{2} \|\text{grad}_S f\|^2$$

where grad_S is the gradient of the surface. Let \mathbf{f} be a vector of a function defined on a surface. The solution for the minimization problem is obtained by solving the Laplace equation with some boundary constraints:

$$\Delta_S \mathbf{f} = \mathbf{0}$$

where Δ_S is the Laplace-Beltrami operator. The consequence of this minimization is that the boundary conditions are smoothly interpolated. When the mappings are harmonic, because the gradient of a mapping is equivalent to the deformation gradient, we are minimizing a local transformation, $\min \|\mathbf{T}\|_F^2$.

There are two important implications that describe relationships between the above mappings. First, every isometric mapping is conformal and equiareal (scale-preserving):

$$\text{isometric} = \text{conformal} \cap \text{equiareal}$$

Second, a conformal mapping is the subspace of harmonic mapping:

$$\text{conformal} \subset \text{harmonic}$$

Thus, conformal mappings are always harmonic but the inverse is not true. Not all harmonic mappings are angle preserving.

From this discussion, we know that for the nonrigid surface registration of different shapes (different scale and details), we should avoid isometric mappings due to their scale-preserving property. Harmonic mappings are also not desirable because they are not angle-preserving (they produce shear distortions). In this paper, we show that conformal mappings are the most appropriate for registering different shapes within the same class.

2.2. Previous work

Nonrigid surface registration techniques Previous surface registration techniques usually adopt one of two types of regularizations: isometric mapping and smooth deformation

(harmonic mapping and its variants). The advantage of the techniques based on isometric mapping [LSP08, HAWG08, TBW*09] is that they can achieve automatic registration. Li et al. [LSP08] achieved isometric registration using the deformation model of [SSP07] that constrains local linear transformations as orthonormal $\mathbf{T}^T \mathbf{T} = \mathbf{I}_d$. Huang et al. [HAWG08] constrain transformations locally as-rigid-as-possible. However, these approaches are incapable of handling models with different sizes or those which undergo large local stretching. In contrast, the techniques based on smoothness regularization—harmonic mapping (membrane model) [WLVP09], bi-harmonic mapping (thin-plate model) [YLSL10] and deformation smoothness [ACP03, SP04, ARV07]—are robust to size differences. However, they are too weak to preserve mesh structures against shear distortions and lose most of template details. Thus, they require many landmarks to obtain a good initial shape.

Conformal mapping In the mesh parameterization field, conformal or angle-preservation mapping is extensively studied and is used to devise efficient low-distortion flattening methods [LPRM02, GY03, DMA02, KSS06, SSP08]. Conformal flattening is also used in shape correspondence algorithms [LF09, KLF11]. However, these techniques require the topology of the template and the target to be the same. Thus, they are not really applicable to real-world 3D scans containing holes and other artifacts. In contrast, our technique is applicable to such scans by constraining transformations to be conformal directly in three-dimensional space. Recently, Paille et al. [PP12] and Martinez et al. [MRT13] proposed the formulations of three-dimensional conformal energies for volumetric parametrization and shape deformation, respectively. Paille et al. [PP12] extended the Cauchy-Riemann equations to 3D. Martinez et al. [MRT13] formulated their deformation framework as continuous shape deformation, which is the generalization of as-killing-as possible deformation [SBCBG11].

The most similar techniques to ours are the methods proposed in [LZW*09] and [PB11], where they achieved registration by constraining deformations as similarity transformations. Liao et al. [LZW*09] use a linear deformation model [SCOL*04]. Thus, it does not handle large rotations. The method in [PB11] is based on shape matching, which starts from a rigidly aligned template and does not incorporate landmark constraints or smoothness regularization (instead, they directly smoothed target positions). Thus, handling large changes in pose or shape and capturing details are difficult for this technique. In contrast, our formulation is based on conformal mapping, where we optimize affine transformations associated with vertices by enforcing the conformal and smoothness constraint directly on the transformations. This allows us to handle large deformations.

3. As-conformal-as-possible surface registration

We propose a surface registration technique that is flexible enough to non-rigidly align two surfaces with different sizes

Table 1: Comparison of mappings. \mathbf{T} is a 3×3 linear transformation matrix. \mathbf{C} is a stretch matrix. \mathbf{I}_d is an identity matrix. s is a scale. \mathbf{R} is a rotation matrix.

Property		$\mathbf{C} = \mathbf{T}^T \mathbf{T}$	\mathbf{T}
Isometric	length-preserving	\mathbf{I}_d	\mathbf{R}
Conformal	angle-preserving	$s^2 \mathbf{I}_d$	$s \mathbf{R}$
Equiareal	scale-preserving		$\det(\mathbf{T}) = 1$
Harmonic	smooth deformation		$\min \ \mathbf{T}\ _F^2$

and details, yet is able to preserve mesh structures of the template. Our goal is to fit the template to scans that share similar semantics but are different in global scale, pose, local scale, and details (two surfaces are in the same class such as different human individuals).

3.1. Notation and overall cost function

The template mesh consists of n vertices and m triangle faces. The vertex positions of the template, $\mathbf{v}_1 \dots \mathbf{v}_n$, are denoted by a $n \times 3$ vector, $\mathbf{v} = [\mathbf{v}_1 \dots \mathbf{v}_n]^T$. The vertex positions of the scan are denoted by $\mathbf{p}_1 \dots \mathbf{p}_N$. The registration is expressed as a set of 3×4 affine transformation matrices $\mathbf{X}_i = [\mathbf{T}_i, \mathbf{t}_i]$ that are associated with each vertex \mathbf{v}_i of the template, where \mathbf{T}_i is a linear transformation and \mathbf{t}_i is a translation. Transformation \mathbf{X}_i is concatenated into a single $4n \times 3$ matrix $\mathbf{X} = [\mathbf{X}_1 \dots \mathbf{X}_n]^T$.

We define the cost function $E(\mathbf{X}) = w_{ACAP} E_{ACAP} + w_{reg} E_{reg} + w_C E_C + w_F E_F$, where E_{ACAP} constrains deformation ACAP, E_{reg} acts as a regularization term to avoid extreme local deformation, E_C penalizes distances between the closest points of template and target surface, and E_F penalizes distances between the feature points of template and target surface.

3.2. Conformal stiffness term

Nonlinear conformal energy Recall that a conformal mapping constrains a local transformation as $\mathbf{T}^T \mathbf{T} = s^2 \mathbf{I}_d$ (Table 1). It must satisfy six conditions: all columns must have the same norm and must be orthogonal to one another.

$$E_{ACAP} = \sum_i \text{Conformal}(\mathbf{T}_i)$$

$$\begin{aligned} \text{Conformal}(\mathbf{T}) &= (\mathbf{c}_1^T \mathbf{c}_2)^2 + (\mathbf{c}_2^T \mathbf{c}_3)^2 + (\mathbf{c}_3^T \mathbf{c}_1)^2 \\ &+ (\mathbf{c}_1^T \mathbf{c}_1 - \mathbf{c}_2^T \mathbf{c}_2)^2 \\ &+ (\mathbf{c}_2^T \mathbf{c}_2 - \mathbf{c}_3^T \mathbf{c}_3)^2 \\ &+ (\mathbf{c}_3^T \mathbf{c}_3 - \mathbf{c}_1^T \mathbf{c}_1)^2 \end{aligned}$$

where \mathbf{c}_1 , \mathbf{c}_2 and \mathbf{c}_3 are the 3×1 column vectors of \mathbf{T} .

Linearization Because E_{ACAP} is nonlinear, it is, in general, expensive to minimize this energy. Following the linear approximation of a linear transformation used in the Laplacian surface editing framework [SCOL*04], we derive the linearized version of E_{ACAP} . In [SCOL*04], a linear trans-




Figure 2: Linear conformal constraint.

formation is approximated as:

$$\mathbf{T} = \begin{pmatrix} s & -h_3 & h_2 \\ h_3 & s & -h_1 \\ -h_2 & h_1 & s \end{pmatrix}$$

We instead enforce a constraint on a transformation as follows:

$$E_{LACAP} = \sum_i \text{LConformal}(\mathbf{T}_i)$$

$$\begin{aligned} \text{LConformal}(\mathbf{T}) &= \|T_{11} - T_{22}\|^2 + \|T_{22} - T_{33}\|^2 + \|T_{33} - T_{11}\|^2 \\ &+ \|T_{12} + T_{21}\|^2 + \|T_{23} + T_{32}\|^2 + \|T_{31} + T_{13}\|^2 \end{aligned}$$

This term constrains the diagonals of \mathbf{T} to be the same and the off-diagonals to satisfy $T_{ij} + T_{ji} = 0$. Figure 2 shows the entries of this constraint in the matrix form. Note that this constraint is equivalent to a 3D extension of the Cauchy-Riemann equation derived by Paille et al. [PP12]; they relate partial derivatives of function, whereas we relate the components of the deformation gradient. This formulation is only valid for a small rotation angle. Thus, we will only use this energy at the last stage of registration, where only small deformations are required.


3.3. Regularization and position constraints

Regularization term We combine two energies for regularization: $E_{reg} = E_{consist} + E_{smooth}$. The role of the first energy is to make the problem well-posed, and the second one avoids extreme local deformation. The first energy makes a linear transformation and translation consistent.

$$E_{consist} = \sum_i \sum_{j \in \mathcal{N}(i)} \|\mathbf{T}_i(\mathbf{v}_j^0 - \mathbf{v}_i^0) + \mathbf{v}_i^0 + \mathbf{t}_i - (\mathbf{v}_j^0 + \mathbf{t}_j)\|^2$$

where $\mathcal{N}(i)$ consists of one-ring neighbors of vertex i . \mathbf{v}_i^0 and \mathbf{v}_j^0 are vertices of the template. The second energy term serves as a regularizer for the deformation by indicating that the linear transformations of adjacent vertices should agree with one another:

$$E_{smooth} = \sum_i \sum_{j \in \mathcal{N}(i)} \|\mathbf{T}_i(\mathbf{v}_j^0 - \mathbf{v}_i^0) + \mathbf{T}_j(\mathbf{v}_i^0 - \mathbf{v}_j^0)\|^2$$



Closest point constraints In order to attract the mesh to the scan, we find the closest point matches of the template and the target surface, $\mathcal{C} = \{(\mathbf{v}_1, \mathbf{p}_{\text{idx}(1)}) \cdots (\mathbf{v}_n, \mathbf{p}_{\text{idx}(n)})\}$, based on the nearest neighbor search, where $\text{idx}(i)$ is the index of the scan point that is matched with vertex i . If the distance between the points exceeds D or if the angle between the normals of the points exceeds Θ , we eliminate that point pair from correspondence set \mathcal{C} .

To avoid extreme deformations in tangential directions, the displacement is projected to the direction of the template normal. Let $\tilde{\mathbf{v}}_i$ and $\tilde{\mathbf{n}}_i$ be the position and normal of template vertex i of the current mesh. Then, displacement is $\alpha_i \tilde{\mathbf{n}}_i$ where $\alpha_i = (\mathbf{p}_{\text{idx}(i)} - \tilde{\mathbf{v}}_i) \cdot \tilde{\mathbf{n}}_i$. The closest point term is defined as:

$$E_C = \sum_{i \in \mathcal{C}} \|\mathbf{v}_i - (\tilde{\mathbf{v}}_i + \alpha_i \tilde{\mathbf{n}}_i)\|^2$$

Feature point constraints The feature point term is established as a soft constraint so that we can control the contribution of this term to the overall energy. The feature correspondences are established automatically or by the user. Suppose n_F feature points are specified. Then the feature point term is defined as follows:

$$E_F = \sum_l^{n_F} \|\mathbf{v}_{\text{idx}(l)} - \mathbf{p}_l\|^2$$

where \mathbf{p}_l is the position of l th feature point and $\text{idx}(l)$ is the index of the corresponding vertex of the template model.

3.4. Optimization

The optimization consists of two loops: The outer loop searches for the nearest neighbor points and constructs the closest point term with w_{reg} set to some value. The inner loop then optimizes the affine transformations at the vertices with the fixed position constraints by minimizing $E(\mathbf{X})$. Once this is converged, w_{reg} is halved and the outer loop finds the closest points again.

Nonlinear least squares We minimize the nonlinear energy $E(\mathbf{X}) = w_{\text{ACAP}}E_{\text{ACAP}} + w_{\text{reg}}E_{\text{reg}} + w_C E_C + w_F E_F$ using an iterative Gauss-Newton method [SSP07]. We unrolled \mathbf{X} and define stacked variables by a $12n \times 1$ column vector \mathbf{x} . The Gauss-Newton algorithm linearizes the nonlinear problem with Taylor expansion about \mathbf{x} :

$$\mathbf{f}(\mathbf{x} + \delta) = \mathbf{f}(\mathbf{x}) + \mathbf{J}\delta$$

The vector $\mathbf{f}(\mathbf{x})$ stacks the squared roots of the cost functions, so that $\mathbf{f}(\mathbf{x})^T \mathbf{f}(\mathbf{x}) = E(\mathbf{x}) = w_{\text{ACAP}}E_{\text{ACAP}} + w_{\text{reg}}E_{\text{reg}} + w_C E_C + w_F E_F$. \mathbf{J} is the Jacobian matrix of $\mathbf{f}(\mathbf{x})$. At each iteration t , we solve a linearized problem and compute an updating vector δ_t to improve the current solution \mathbf{x}_t :

$$\begin{aligned} \mathbf{J}_t^T \mathbf{J}_t \delta_t &= -\mathbf{J}_t^T \mathbf{f}(\mathbf{x}_t) \\ \mathbf{x}_{t+1} &= \mathbf{x}_t + \delta_t \end{aligned} \quad (1)$$

In each Gauss-Newton iteration, we solve the normal equations by Cholesky factorization. We must calculate both the

symbolic and numeric factorization of \mathbf{J}_t once after the outer loop finds the closest points. In the inner loop, however, the non-zero structure of \mathbf{J}_t remains unchanged. Thus, we can reuse the symbolic factorization to speed up computations. The inner loop typically takes 6 iterations until convergence.

Linear least squares The linearized conformal registration energy $E(\mathbf{x}) = w_{\text{ACAP}}E_{\text{LACAP}} + w_{\text{reg}}E_{\text{reg}} + w_C E_C + w_F E_F$ is minimized in a linear least squares sense as:

$$\mathbf{A}^T \mathbf{A} \mathbf{x} = \mathbf{A}^T \mathbf{b} \quad (2)$$

where the system matrix \mathbf{A} and the right hand side vector \mathbf{b} are defined from the constraints.

3.5. Relation to as-similar-as-possible energy

As we can see from Table 1, we can also achieve conformal mapping by constraining local linear transformations to similarity transformations, i.e., minimizing the as-similar-as-possible (ASAP) energy:

$$E_{\text{ASAP}} = \sum_i \|\mathbf{T}_i - \tilde{s}_i \tilde{\mathbf{R}}_i\|_F^2$$

where scale \tilde{s}_i and rotation $\tilde{\mathbf{R}}_i$ are calculated from the current transformations. Within our framework, this energy can be minimized by alternating between the optimization of affine transformations and the calculations of similarity transformations using polar decompositions.

Note that E_{ACAP} and E_{ASAP} are basically equivalent. However, because they employ different optimization techniques, differences appear in deformation quality, convergence and single-iteration time. We will compare ACAP and ASAP from these aspects in Section 6.

4. Coarse-to-fine fitting strategy

To achieve registration efficiently and robustly, we take a coarse-to-fine fitting strategy. This not only improves performance but also reduces risks of generating fold-overs. We use two slightly different techniques. The first strategy fits a coarse template to the scan, performs subdivision and then fits the resulting mesh to the scan again. The other strategy incorporates space deformation and performs an efficient subspace deformation technique by layering a coarse graph under the dense mesh. Once subspace deformation is done, we perform registration on the original dense resolution.

4.1. Fitting Steps

The fitting steps are summarized as follows:

Step1: Initial fitting The template is first roughly fitted to several feature points. We require at least three correspondences to achieve the initial fitting. At this stage, affine transformations are associated to vertices of a coarse mesh (graph) and the overall size of the template is adjusted. Here, we use the nonlinear conformal constraint to handle large rotations.




Figure 3: Effect of using different distance functions and K values during subspace deformation. The use of Euclid distances leads to distortions, where geodesically far apart regions are wrongly deformed. Diffusion distances can avoid this problem. $K = 8$ balances quality and efficiency.

Step2: Mid-scale fitting After fitting the template roughly to the scan using feature points, the mesh is deformed gradually toward the scan. Again, affine transformations are associated to vertices of the coarse mesh (graph) and optimized nonlinearly. At this stage, we focus on adjusting local scale and we do not aim to capture details.

Step3: Subdivision (optional)

Step4: Fine fitting In this stage, the focus is on capturing details. Thus, affine transformations are associated to vertices of the dense (original or subdivided) mesh. Here, to improve efficiency, we minimize the linear conformal registration energy.

4.2. Subdivision

The strategy based on subdivision starts from a coarse mesh with several hundred vertices. We first align the coarse template and the scan using ACAP registration. We then subdivide the deformed mesh to generate a dense result. Any subdivision method can be used here. We use Loops subdivision technique [Loo87]. Once we obtain a dense mesh, we perform registration again to attract the vertices of the dense mesh to the scan surface.

4.3. Subspace deformation

To align a dense template efficiently toward a scan, we use a subspace deformation technique called embedded deformation [SSP07]. The embedded deformation technique layers coarse graphs under the dense mesh and solves the problem in the reduced space, which can significantly speed up computations. An affine transformation is associated with each vertex in the coarse graph. Each vertex in the dense mesh is assigned skinning weights and its deformed position is approximated from transformations of coarse graphs.

The vertex position of the coarse graph is represented by \mathbf{v}_k . The linear transformation and the translation associated with \mathbf{v}_k is denoted by \mathbf{T}_k and \mathbf{t}_k , respectively. Let \mathbf{v}_k^0 be the vertex position of the coarse mesh in the rest state. The deformed vertex $\bar{\mathbf{v}}_i$ is obtained as follows:

$$\bar{\mathbf{v}}_i = \sum_{k=1}^c w_i^{(k)} [\mathbf{T}_k(\mathbf{v}_i^0 - \mathbf{v}_k^0) + \mathbf{v}_k^0 + \mathbf{t}_k]$$

c is the number of vertices in the coarse graph. $w_i^{(k)}$ is a weight for vertex i , controlling how much \mathbf{v}_i is influenced by \mathbf{v}_k .

The coarse graph can be established using mesh simplification techniques or the farthest point sampling strategy. The spatially varying weight $w_i^{(k)}$ for each vertex is computed by a K -nearest approach as:

$$w_i^{(k)} = (1 - d_{ik}/d_{max})^2$$

and then normalized in order to sum to one. Here, d_{ik} is the distance from point i to point k and d_{max} is the distance to the $K + 1$ -nearest node. We empirically determined $K = 8$. There are several choices for distance function, such as Euclid, geodesic and diffusion. Euclid distances are fast and easy to compute, but there is the risk of connecting the geodesically far apart regions. Geodesic distances can prevent this problem but their computational cost is high. Diffusion distances mimic geodesics and are relatively easy to compute. An additional benefit of using diffusion distances is that they are robust to topological errors, which means that they are applicable to surfaces with holes and topological short circuits. We use Euclid distances for the face model, as it is rarely the case that geodesically far apart regions are close in Euclidean space. As for the body model, we use diffusion distance [dGGV08]. Figure 3 shows the benefit of using diffusion distances and $K = 8$ during subspace deformation.

Now the feature point constraints are enforced as:

$$E_F = \sum_l^{n_F} \|\bar{\mathbf{v}}_{\text{id}_x(l)} - \mathbf{p}_l\|^2$$

With this, the user can specify landmarks anywhere freely on the dense mesh.

4.4. Weights and parameters

In every stage, $D = 0.02r_{\text{box}}$, where r_{box} is the bounding box diagonal, and $\Theta = 90^\circ$ are used. In the initial fitting stage, the closest point term is ignored and a large value for the weight of the landmark term is chosen: $w_{\text{ACAP}} = 1000$, $w_{\text{reg}} = 1000$, $w_C = 0$ and $w_F = 10^5$. In the mid-scale fitting




Figure 4: We provide a simple user interface (UI). The user first selects several (at least three) correspondences. The system outputs an initial fitting result. The user then adjusts stiffness using a slider bar (mid-scale fitting). The user will then add and remove correspondences by observing the result (feature refinement). If the semantic correspondences are satisfactory, the process is terminated. Finally, the registration is performed off-line on the dense resolution.

stage, we set $w_C = 10$, and w_{reg} is halved until w_{reg} reaches 1. In the fine fitting stage, we took the same procedure with $w_F = 1$.

5. Applications

Based on ACAP surface registration, we propose two use-case scenarios for surface registration, which reduces user interventions during template fitting.

5.1. Automatic registration

The first step of the automatic registration is to establish point-to-point correspondences between two surfaces. In principle, any matching technique can be employed. However, in our case, the shape and size of the target surface differs from that of the template. The target scan also contains holes. Thus, we use a non-rigid shape matching technique that is robust against data incompleteness and non-isometric deformations [YYYS14].

Once correspondences are established (100–300 correspondences are obtained), we feed them into the ACAP registration framework as feature point constraints to achieve initial fitting. We next iteratively deform the mesh toward the target by incorporating the closest point constraints. The above technique allows us to automatically align the template to scans in various shapes and poses.

5.2. Interactive registration

Although our automatic registration technique can be performed without any user interactions, in practice it is often the case that fully automatic matching algorithms output low-quality correspondences or miss matches. Therefore, user specification is still the most reliable way to provide correspondences. Thus, designing a means for reducing the effort required for this task is very important. This is also beneficial for the modification of low-quality correspondences produced by automatic methods. We therefore design an interactive user interface to help users specify the correspondences between the template and target shapes in

Table 2: Statistics for the examined models.

Template	Scan	#CV	#DV	#SV	#L	Figure	
Head1	Hair	537	15k	100k	5	Fig. 1	
Head1	Ear	537	15k	50k	7	Fig. 4	
Head1	Female	537	15k	15k	5	Fig. 5	
Head2	Male	570	10k	100k	4	Fig. 5	
Head1	Laugh	537	15k	14k	16	Fig. 5	
SCAPE	Abhijeet	500	12k	80k	5	Fig. 6	
	Ilya	Abhijeet	500	20k	80k	–	Fig. 7
SCAPE	SCAPE	500	12k	65k	–	Fig. 7	
SCAPE	Gorilla	500	–	15k	6	Fig. 12	
Horse	Camel	1000	–	20k	12	Fig. 13	

#CV and #DV indicate the number of vertices on a coarse and dense mesh, respectively. #SV indicates the number of the scans vertices. #L is the number of landmarks.

an intuitive manner. The proposed interactive system helps users freely add, remove and modify correspondences while observing the result on-the-fly. The user can also change the stiffness interactively with slider bars.

A typical workflow is illustrated in Figure 4. The user first selects several (at least three) correspondences. The system outputs an initial fitting result. The user then adjust stiffness using a slider bar (mid-scale fitting). The user will then add and remove correspondences by observing the result (feature refinement). This process repeats until the user is satisfied with the overall registration. Finally, the registration is performed off-line on the dense resolution (Fine fitting).

During the interactive registration, we only optimize transformations of the coarse graph. Also, the closet points are searched only from the vertices of the coarse graph. After interactive registration, we perform surface registration off-line on the dense mesh. These strategy allows for registration of a large mesh. Note that, with the use of subspace deformation, the user can specify feature point constraints freely on the dense mesh and move them interactively.

6. Experiments

We tested our method on a wide variety of surfaces, i.e., face and whole-body models with different expressions and poses. Most of the scans contain many holes, occluded regions and noise. Statistics of the models are shown in Table 2. We compare our method with other state-of-the-art methods qualitatively and quantitatively.

Wide range of models Figures 1, 5, 6 and 9 show the results of interactive registration. Figure 7 shows the result of automatic registration. Given a smooth template, the proposed technique is able to capture geometrical details such as the ones exhibit around the hairs and ears (Figs 1 and 9). Notice how the connectivity is preserved nicely after large deformation (Fig. 1). Furthermore, attributing to the conformal stiffness that automatically adjusts local scale, occluded regions and holes are filled in a visually pleasing manner. With our registration technique, large deformations due to




Figure 7: Automatic registration to scans with holes. Correspondences are served as position constraints of initial fitting. We then iteratively deform the mesh towards the target by incorporating the closest point constraint. Note that the automatic registration technique is not always successful (see the right figure where the ear and nose are deformed to semantically incorrect places due to failed correspondences).

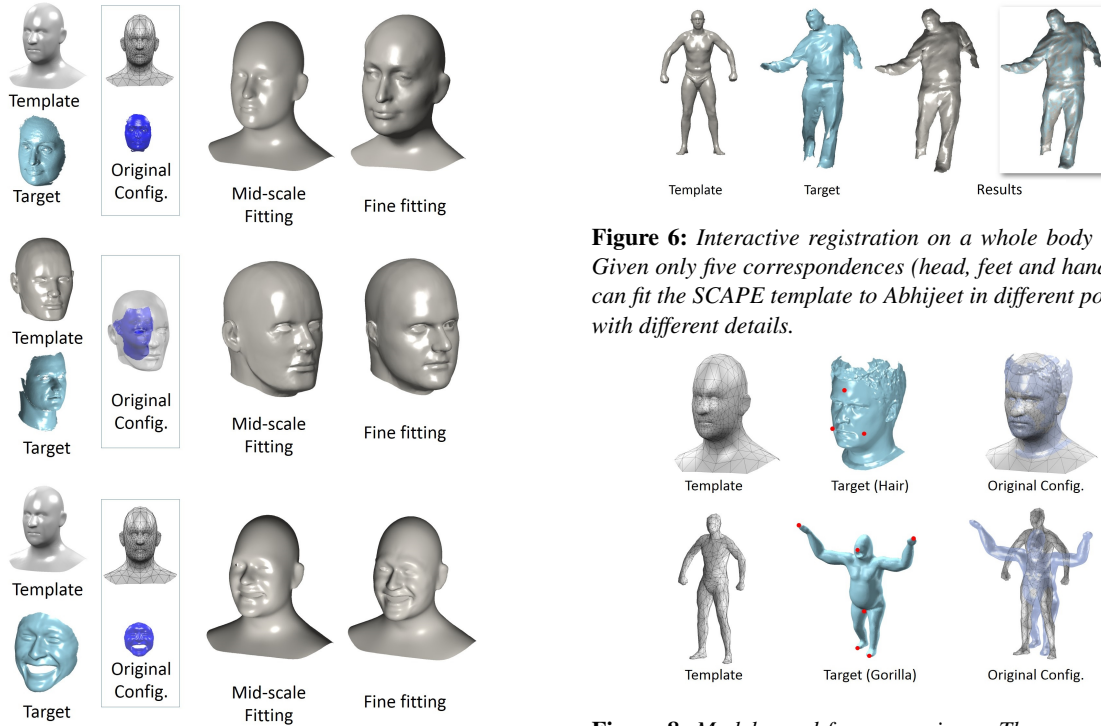


Figure 5: Results on face scans using our interactive registration.

joint extensions or flexions can be handled without generating low-quality triangles (Figs. 6 and 7).

Comparisons We compare our technique (ACAP, ASAP and LACAP) with other state-of-the-art algorithms: the embedded deformation technique (ED) [SSP07] that is used by Li et al. [LSP08], the linear Laplace deformation technique based on implicit optimization (ILD) [SCOL*04] which is used in [LZW*09], the shape-matching based registration technique that minimizes the as-similar-as-possible energy (SM-ASAP) [PB11], the registration technique that uses the triangle-based deformation smoothness regularization (TDS) [SP04] and the registration technique that uses the point-based deformation smoothness regularization (PDS) [ARV07]. ED is an isometric counterpart of ACAP, in which,

Figure 6: Interactive registration on a whole body model. Given only five correspondences (head, feet and hands), we can fit the SCAPE template to Abhijeet in different pose and with different details.

Figure 8: Models used for comparisons. Three correspondences are specified for head registration (Hair). Six correspondences are specified for whole-body registration (Gorilla).

for this comparison, we replaced the nonlinear conformal stiffness with the isometric one and used the regularization term presented in this paper. ILD and SM-ASAP attempt to minimize as-similar-as-possible (conformal) energies. SM-ASAP uses shape matching, which starts from rigidly aligning the template to the target. It does not require specifying feature points. TDS and PDS are based on smoothness regularization.

The models used for comparisons are shown in Fig. 8. Note that global scale, orientation and position of the models are pre-aligned. We measure 1) data error, the average distance from the vertices of the deformed template to the corresponding points of the scan relative to the bounding box




Figure 9: Linear vs nonlinear ACAP on fine fitting.

diagonal, 2) angle error, the average angle deviation from the template and 3) bending error, the average deviation in dihedral angles.

Figure 12 shows qualitative comparisons. For the Hair model, which requires relatively small deformations for registration, ACAP, ASAP, LACAP and ED [SSP07] produce visually similar results. In contrast, ILD [SCOL*04], SM-ASAP [PB11], TDS [SP04] and PDS [ARV07] produce the results with shear distortions and fold-overs. For the Gorilla model, the results of ACAP and ASAP exhibit almost no distortions, whereas others reveal fold-overs and distortions. LACAP, although linear, produces comparable results until registration requires large rotations; e.g., fold-overs occurred around the arms when fitting to Gorilla in a different arm pose. At a glance, ED [SSP07] generates, visually similar results to ACAP but the feet of Gorilla exhibits fold-overs. A possible explanation for this was that ED could not adjust the scale differences between the legs of Gorilla and Human (Fig 13a). Thus, the heels were deeply nailed into the ground and the vertices around the ankles were then attracted to the heels. The ability of ACAP to adjust local scale can be clearly seen when fitting the Horse into the Camel where the size of the tail is adjusted properly (Fig. 13c). SM-ASAP [PB11] is quite fast as it does not require solving linear systems, but it cannot handle large changes in pose and shape. ILD [SCOL*04] cannot handle large rotations. We also found that ILD becomes unstable when the models have boundaries and it requires users to specify additional features there. PDS [ARV07] is prone to shear distortions and it tends not to work well with feature points that are placed in a coplanar arrangement. TDS [SP04] generates an initial shape with extreme shrinkage, which results in large fold-overs in the final result. Table 3 shows quantitative comparisons. The angle error of ACAP is the smallest among all the techniques. In addition, the bending error on the Gorilla example is the lowest of all, which reflects the ability of ACAP to reduce the risk of producing fold-overs.

In Fig. 9, we also compare fine-fitting results obtained using the linear and nonlinear conformal constraints (LACAP vs ACAP). The results look very similar, almost without visually noticeable differences. In fact, the difference of the angle error from mid-scale fitting to fine fitting is only 0.15° . We therefore believe that LACAP is sufficiently accurate when used in fine fitting.

Performance We implemented the prototype of our algo-




Figure 10: Convergences of ACAP and ASAP are compared. The conformal energy E_{ACAP} is measured.

rithm in Matlab, partially written in C/C++ on an Intel Core i7 3.4GHz 64-bit workstation. CHOLMOD [CDHR08] is used for constructing, factorizing and solving Eqs. (1) and (2). The timing is shown on Table 4. The reuse of the symbolic factorization reduces the time required for a single Gauss-Newton iteration by about 35%. The linearized version of ACAP is approximately five times faster than the nonlinear version for the problem size of fine fitting.

Number of feature points required from users The numbers of feature points required from the user are shown on Table 2. Previous techniques [SP04, YLSL10] require specifying 20–70 feature points, whereas our technique requires less than 20 points. This is because ACAP provides a good initial shape from a few feature points and can avoid extreme distortions during iterative fitting.

ACAP vs ASAP ASAP is an alternative formulation of ACAP, but the quality of its results is slightly lower than that of ACAP (Table 3). On the other hand, ASAP is approximately twice as fast as ACAP because the system matrix of ASAP stays in place during a single outer loop, until the closest point term changes. However, as shown in Fig. 10, ASAP's convergence is extremely slower than ACAP. Consequently, ACAP is preferred over ASAP, especially when registration requires large deformations.

Shape interpolation Because the models obtained using our technique have identical connectivities, we can interpolate the models and obtain shapes in-between. We used a shape interpolation technique based on mass-springs systems similar to the method proposed in [MWF*12]. The results are shown in Fig. 11.

Limitations One limitation of our algorithm is the performance due to the use of nonlinear optimization (nonlinear ACAP). Although our technique can achieve near interactive rates by making careful design choices (analytically building Jacobian, reusing symbolic factorization, coarse-to-fine strategy), it is slower than previous techniques; according to performance comparisons on mid-scale fitting based on our implementations, the previous techniques proposed in [PB11] and [LZW*09] are approximately 10 and 3 times faster than ACAP, respectively. The performance

Table 3: Quantitative comparisons. *D, A and B indicate data error [%], angle error [deg] and bending error[deg], respectively.*

	ACAP			ASAP			LACAP			ED			SM-ASAP			ILD			PDS			TDS		
	D	A	B	D	A	B	D	A	B	D	A	B	D	A	B	D	A	B	D	A	B	D	A	B
Hair	0.48	3.3	12.0	0.47	3.6	11.4	0.45	5.3	13.3	0.48	4.4	11.7	1.17	8.9	13.7	0.58	4.8	15.7	0.95	13.9	14.7	1.61	23.9	26.5
Ear	0.40	2.9	5.8	0.44	3.5	5.6	0.42	4.6	7.7	0.40	3.5	6.9	0.67	10.1	13.7	0.45	13.3	13.3	0.43	7.3	11.7	0.89	18.7	24.6
Gorilla	0.25	10.7	14.3	0.31	12.3	15.9	0.31	11.9	20.4	0.26	11.8	21.1	4.76	11.5	17.2	3.23	13.8	23.6	0.37	20.5	27.5	0.41	12.1	20.4

Table 4: Timings(in seconds).

	Meshes			Mid-scale fitting			Fine fitting		
	#Coarse	#Dense	#Target	NN search	GN	Regist.	NN search	LS	Regist.
Subdivision	537	15k(100k)	50k	0.020	0.022	0.25	0.08 (0.376)	0.1 (0.8)	4.5 (43)
Subspace	570	10k	100k	0.050	0.023	0.26	0.13	0.07	2.7

#Coarse, #Dense and #Target indicate the mesh size. “NN search,” “GN,” “LS,” and “Regist.” indicate the time required for the nearest neighbor search, a single Gauss Newton iteration, a single least squares solve and a single iteration of ACAP registration, respectively.





Figure 11: Shape interpolation from the SCAPE model to Abhijeet.

of our method might be improved by using the linearized version (LACAP) and having the user specify rotation constraints to handles.

Another limitation is fold-overs. Although ACAP with the coarse-to-fine strategy is quite robust against fold-overs, it does not guarantee to solve this issue, especially for models with a complex topology. Thus, the fold-over removing technique [YLSL10] or bounded distortion mapping [Lip12] would be useful for solving this problem thoroughly. An easy way to address the problem is to add one or two landmarks around the fold-overs and move the landmarks slightly in order to stretch the region.



7. Conclusion

We have presented a novel surface registration method that constrains deformations locally as-conformal-as possible. With the proposed conformal stiffness term, mesh structures and the quality of the template are preserved by keeping the original angles of triangle meshes. We believe that both automatic and interactive registration based on ACAP technique can help reduce user efforts required during template fitting. Experimental results showed that our technique preserves angles better than previous techniques and can reduce the risk of producing fold-overs.

In the future, we would like to find a way to fuse the interactive and automatic registration in order to compromise each other. With the interactive system, the user can specify correspondences relatively accurately anywhere on a surface; however, sometimes this is tedious (e.g. 10 landmarks were needed to align the lips of the Laugh model). On the other hand, automatic registration does not require user interaction, but cannot produce accurate correspondences, especially for the regions where stable features are not available. We believe that the fusion of these two techniques would lead to a better tool for surface registration.

Acknowledgements

We thank anonymous reviewers for their valuable comments. We thank Drago Angelov, Ilya Baran, Bob Sumner and Daniel Vlasic for their 3D models. We thank Graphics Lab at Institute for Creative Technologies, University of Southern California to provide some of the facial scans, and J. P. Lewis for his valuable comments. This work was partially supported by JSPS KAKENHI Grant Number 25880033.

References

[ACP03] ALLEN B., CURLESS B., POPOVIĆ Z.: The space of human body shapes: reconstruction and parameterization from range scans. *ACM Trans. Graph.* 22, 3 (2003), 587–594.

[ARV07] AMBERG B., ROMDHANI S., VETTER T.: Optimal Step Non-rigid ICP Algorithms for Surface Registration. In *CVPR* (2007).

[CDHR08] CHEN Y., DAVIS T. A., HAGER W. W., RAJAMANICKAM S.: Algorithm 887: CHOLMOD, Supernodal Sparse Cholesky Factorization and Update/Downdate. *ACM Trans. Math. Softw.* 35 (2008).

[dGGV08] DE GOES F., GOLDENSTEIN S., VELHO L.: A hierarchical segmentation of articulated bodies. In *Proceedings of the Symposium on Geometry Processing* (2008), pp. 1349–1356.

[DMA02] DESBRUN M., MEYER M., ALLIEZ P.: Intrinsic Parameterizations of Surface Meshes. *Comput. Graph. Forum* 21, 3 (2002), 209–218.

[FH05] FLOATER M. S., HORMANN K.: *Surface Parameterization: a Tutorial and Survey*, 2005.

[GY03] GU X., YAU S.-T.: Global conformal surface parameterization. In *Proceedings of the 2003 Eurographics/ACM SIGGRAPH symposium on Geometry processing* (2003), SGP '03, pp. 127–137.

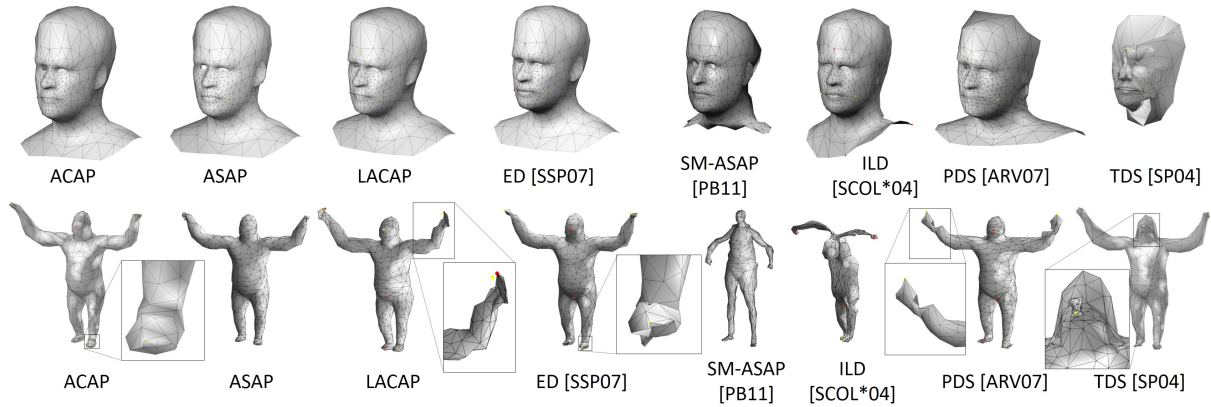


Figure 12: Comparisons. The experimental setup is shown in Fig. 8.




Figure 13: ACAP vs ED [SSP07]. a) Gorilla, b) Hair. The intermediate results of Figure 12 with compatible w_{reg} values are shown. ED does not adjust local scale, which could lead to fold-overs, shear distortions, and slow convergence; c) Horse to Camel: The ability of ACAP to adjust local scale can be clearly seen from around the tail.

[HAWG08] HUANG Q.-X., ADAMS B., WICKE M., GUIBAS L. J.: Non-rigid registration under isometric deformations. In *Proceedings of the Symposium on Geometry Processing* (2008), pp. 1449–1457.

[KLF11] KIM V. G., LIPMAN Y., FUNKHOUSER T.: Blended intrinsic maps. *ACM Trans. Graph.* 30, 4 (2011), 79:1–79:12.

[KSS06] KHAREVYCH L., SPRINGBORN B., SCHRÖDER P.: Discrete conformal mappings via circle patterns. *ACM Trans. Graph.* 25, 2 (2006), 412–438.

[LF09] LIPMAN Y., FUNKHOUSER T.: Mobius voting for surface correspondence. *ACM Trans. Graph.* 28, 3 (2009), 72:1–72:12.

[Lip12] LIPMAN Y.: Bounded Distortion Mapping Spaces for Triangular Meshes. *ACM Trans. Graph.* 31, 4 (July 2012), 108:1–108:13.

[Loo87] LOOP C.: *Smooth subdivision surfaces based on triangles.* Master’s thesis, 1987.

[LPRM02] LÉVY B., PETITJEAN S., RAY N., MAILLOT J.: Least squares conformal maps for automatic texture atlas generation. *ACM Trans. Graph.* 21, 3 (2002), 362–371.

[LSP08] LI H., SUMNER R. W., PAULY M.: Global correspondence optimization for non-rigid registration of depth scans. In *Proceedings of the Symposium on Geometry Processing* (2008), pp. 1421–1430.

[LZW*09] LIAO M., ZHANG Q., WANG H., YANG R., GONG M.: Modeling deformable objects from a single depth camera. In *ICCV* (2009), pp. 167–174.

[MRT13] MARTINEZ ESTURO J., RÖSSL C., THEISEL H.: Generalized Metric Energies for Continuous Shape Deformation. *Springer LNCS (Proc. Curves and Surfaces 2012) 8177*, 1 (2013), 135–157.

[MWF*12] MA W.-C., WANG Y.-H. T., FYFFE G., CHEN B.-Y., DEBEVEC P. E.: A blendshape model that incorporates physical interaction. *Journal of Visualization and Computer Animation* 23, 3-4 (2012), 235–243.

[PB11] PAPAJOV C., BURSCHKA D.: Deformable 3D Shape Registration Based on Local Similarity Transforms. *Computer Graphics Forum* 30 (2011).

[PP12] PAILLÉ G.-P., POULIN P.: As-conformal-as-possible discrete volumetric mapping. *Computers & Graphics* 36, 5 (2012), 427–433.

[SBCBG11] SOLOMON J., BEN-CHEN M., BUTSCHER A., GUIBAS L. J.: As-Killing-As-Possible Vector Fields for Planar Deformation. *Comput. Graph. Forum* 30, 5 (2011), 1543–1552.

[SCOL*04] SORKINE O., COHEN-OR D., LIPMAN Y., ALEXA M., RÖSSL C., SEIDEL H.-P.: Laplacian surface editing. In *Proc. of Symposium on Geometry Processing* (2004), pp. 175–184.

[SP04] SUMNER R. W., POPOVIĆ J.: Deformation transfer for triangle meshes. *ACM Trans. Graph.* 23 (2004), 399–405.

[SSP07] SUMNER R. W., SCHMID J., PAULY M.: Embedded deformation for shape manipulation. *ACM Trans. Graph.* 26 (2007).

[SSP08] SPRINGBORN B., SCHRÖDER P., PINKALL U.: Conformal Equivalence of Triangle Meshes. *ACM Trans. Graph.* 27, 3 (2008), 77:1–77:11.

[TBW*09] TEVS A., BOKELOH M., WAND M., SCHILLING A., SEIDEL H.-P.: Isometric registration of ambiguous and partial data. In *CVPR* (2009), pp. 1185–1192.

[WLVP09] WEISE T., LI H., VAN GOOL L., PAULY M.: Face/Off: live facial puppetry. In *Proceedings of the 2009 ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (2009), pp. 7–16.

[YLSL10] YEH I.-C., LIN C.-H., SORKINE O., LEE T.-Y.: Template-based 3D Model Fitting Using Dual-domain Relaxation. *IEEE Transactions on Visualization and Computer Graphics* 99 (2010).

[YYYS14] YOSHIYASU Y., YOSHIDA E., YOKOI K., SAGAWA R.: Symmetry-Aware Nonrigid Maching of Incomplete 3D Surfaces. In *Proceedings of Computer Vision and Pattern Recognition (CVPR)* (2014).