Robust Nonrigid Registration by Convex Optimization

Qifeng Chen — Stanford University
Vladlen Koltun — Intel Labs

Abstract

We present an approach to nonrigid registration of 3D surfaces. We cast isometric embedding as MRF optimization and apply efficient global optimization algorithms based on linear programming relaxations. The Markov random field perspective suggests a natural connection with robust statistics and motivates robust forms of the intrinsic distortion functional. Our approach outperforms a large body of prior work by a significant margin, increasing registration precision on real data by a factor of 3.

Preliminaries

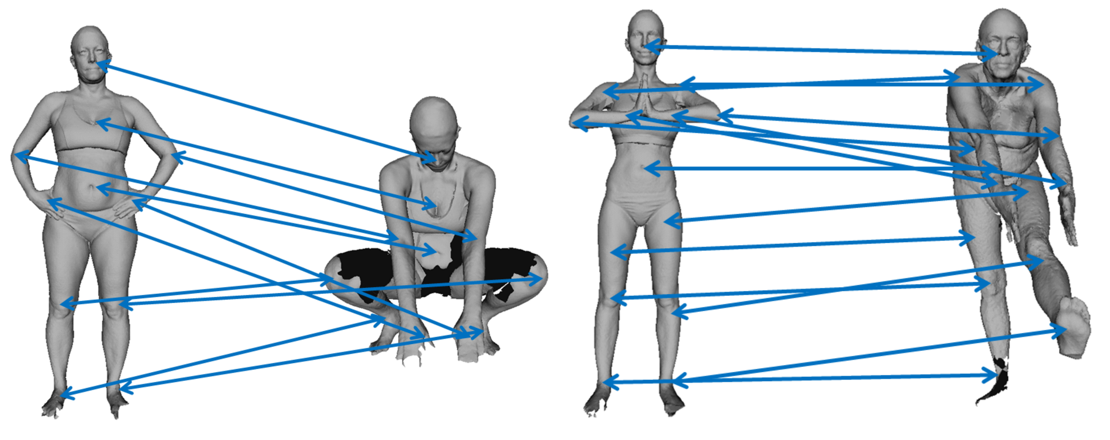

intra-subject registration — inter-subject registration

Figure 1: Nonrigid registration of 3D surfaces.

Given two surfaces $\mathcal{S}, \mathcal{T} \subset \mathbb{R}^3$, the nonrigid registration problem calls for computing a physically or perceptually meaningful mapping $f : \mathcal{S} \to \mathcal{T}$.
A direct approach to this problem is to minimize the intrinsic distortion induced by the mapping [2]:

$$E(f) = \sum_{i,j} w_{i,j} \left| d_{\mathcal{S}}(s_i, s_j) - d_{\mathcal{T}}(f(s_i), f(s_j)) \right|^p. \qquad (1)$$

Here $d_{\mathcal{S}}(\cdot,\cdot)$ and $d_{\mathcal{T}}(\cdot,\cdot)$ are the geodesic distance functions on $\mathcal{S}$ and $\mathcal{T}$. We show that global optimization of a variant of this classical objective yields state-of-the-art registration accuracy, outperforming far more complex registration pipelines. In particular, we do not use keypoints, segments, or shape descriptors. Our work demonstrates that a direct approach that simply optimizes an isometric distortion objective substantially outperforms descriptor-based techniques.

Optimization

The embedding objective (1) corresponds to a continuous Markov random field with pairwise potentials. To apply optimization techniques developed for discrete Markov random fields, we discretize the label space. The discrete MRF objective is

$$\underset{\mathbf{l}}{\text{minimize}} \sum_{i,j} w_{i,j} \left| d_{\mathcal{S}}(s_i, s_j) - d_{\mathcal{T}}(t_{l_i}, t_{l_j}) \right|^p. \qquad (2)$$

Each labeling $\mathbf{l}$ specifies a correspondence between the sample sets S and T.
The MRF optimization problem is NP-hard in general. A natural approach to deriving approximate algorithms is to represent the problem as an integer linear program and relax the integer constraints to obtain a linear program. The LP relaxation is

$$\begin{aligned} \underset{\mathbf{0} \le \mathbf{x} \le \mathbf{1}}{\text{minimize}} \quad & \sum_{i,j} \sum_{a,b} \theta_{ij}(a,b) x_{ij}^{ab} \\ \text{s.t.} \quad & \sum_b x_{ij}^{ab} = x_i^a, \quad \forall i,j \in S,\ \forall a \in L \\ & \sum_a x_i^a = 1, \qquad \forall i \in S \\ & \theta_{ij}(l_i, l_j) = w_{i,j} \left| d_{\mathcal{S}}(s_i, s_j) - d_{\mathcal{T}}(t_{l_i}, t_{l_j}) \right|^p. \end{aligned} \qquad (3)$$

Here $\{x_i^a\}$ and $\{x_{ij}^{ab}\}$ are auxiliary variables that specify a distribution over the space of labelings. The first set of constraints imply that correspondences between pairs of samples must be consistent with correspondences between individual samples.
While standard LP solvers can be used to solve (3), they do not take advantage of the structure of the problem and do not scale well. Over the past decade, a variety of special-purpose MRF optimization algorithms have been developed that can effectively optimize large models. Many of the most successful ones can be interpreted in terms of the dual of (3). The dual LP is

$$\begin{aligned} \underset{\boldsymbol{\delta}}{\text{maximize}} \quad & \sum_i \min_a \tilde{\theta}_i^a + \sum_{i,j} \min_{a,b} \tilde{\theta}_{ij}^{ab}, \\ \text{where} \quad & \tilde{\theta}_i^a = \sum_j \delta_{ji}^a \\ & \tilde{\theta}_{ij}^{ab} = \theta_{ij}(a,b) - \delta_{ji}^a - \delta_{ij}^b \end{aligned} \qquad (4)$$

The variables $\{\delta_{ij}^a\}$ have a natural interpretation as messages. Optimization of the dual LP yields revised potentials $\tilde{\boldsymbol{\theta}}$, which can be used to obtain an optimized configuration. In particular, we use sequential tree-reweighted message passing (TRW-S) [4] to optimize objective (2) and its variants, as TRW-S exhibits both high accuracy and rapid convergence.

Objective

To improve the objective, we replace the unbounded L^p norm with a robust truncated L^1 norm. Furthermore, we additionally apply a Laplace weight to attenuate the contribution of long-distance pairs to the objective as large geodesic distances usually trigger high penalties. Our combined penalty function is

$$\rho(x,y) = \exp\left(-\frac{\min(x,y)}{b} \right) \min\left(|x-y|, \tau \right),$$

We also add a weak extrinsic unary term into the objective to disambiguate intrinsic symmetries. The complete objective is

$$E(\mathbf{l}) = \sum_{i,j=1}^{n} \rho\left(d_{\mathcal{S}}(s_i, s_j), d_{\mathcal{T}}(t_{l_i}, t_{l_j}) \right) + \lambda \sum_i \| s_i - t_{l_i} \|. \qquad (5)$$

Implementation

Preprocessing. Our approach does not assume that the input surfaces are watertight. To simplify the computation of geodesic distances, we begin by applying Poisson reconstruction to the meshes. Then we sample points on the watertight surfaces using farthest point sampling.

Global optimization. Our objective can be used to globally optimize the registration objective for hundreds of samples on each surface on a typical workstation. Since the primal solution produced by TRW-S is related to the ordering of nodes, we permute the ordering 10 times and choose the best primal solution.

Upsampling and refinement. If desired, we can upsample the solution computed by global optimization to thousands of correspondences, which can be further refined by fusion moves.

Evaluation

Results are summarized in Figures 3 and 4. Our approach outperforms a large body of prior work by a multiplicative factor.

Blended intrinsic maps [3] — Random forests [5] — Our approach

Figure 2: Nonrigid registration of scans from the FAUST dataset [1].

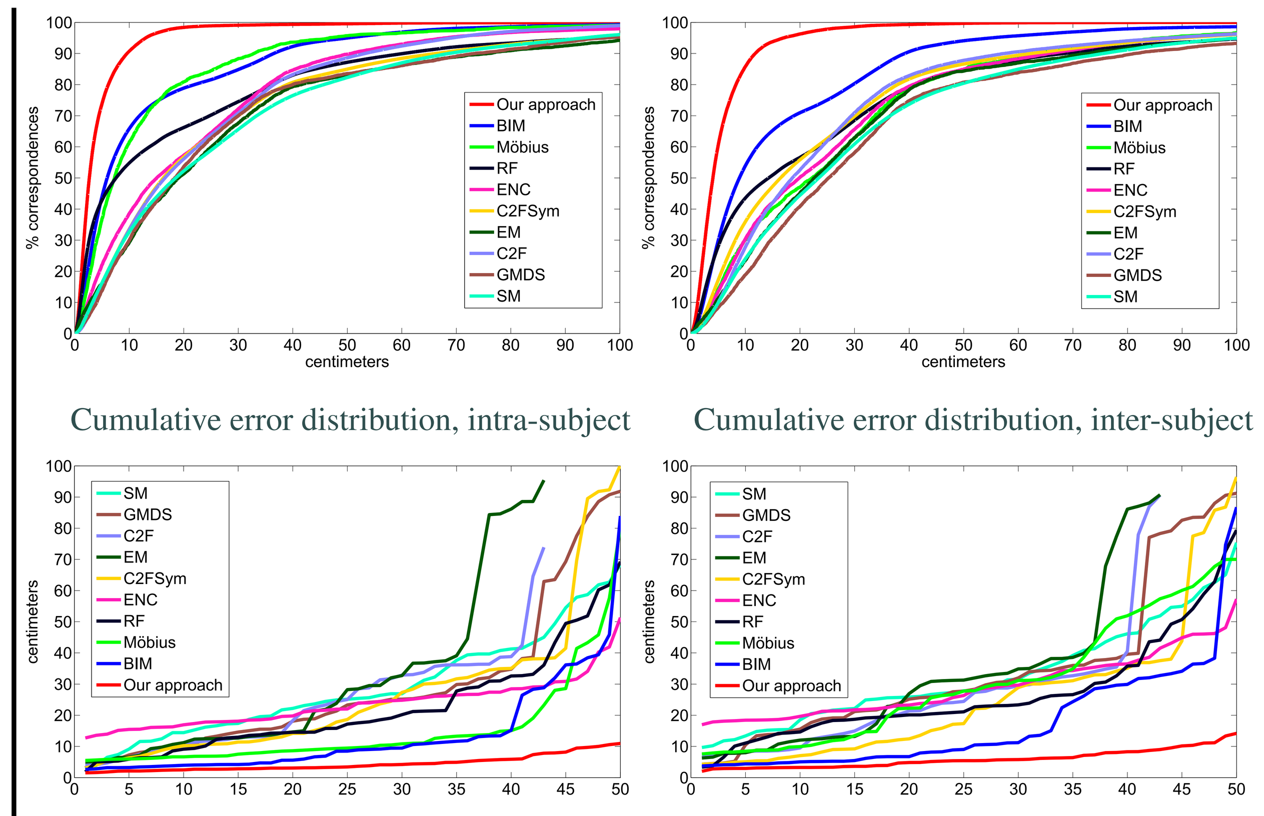

Cumulative error distribution, intra-subject — Cumulative error distribution, inter-subject

Average error for each intra-subject pair — Average error for each inter-subject pair

Figure 3: Evaluation on the FAUST dataset. Top: Cumulative error distributions, in centimeters (higher is better). Bottom: Average error for each of the tested pairs, sorted by magnitude (lower is better).

	intra-subject			inter-subject		
method	AE (cm)	worst AE	10cm-recall	AE (cm)	worst AE	10cm-recall
GMDS [2]	28.98	91.84	0.300	35.06	91.21	0.188
Möbius	14.99	80.40	0.614	30.58	70.02	0.300
BIM [3]	13.60	83.90	0.658	17.36	86.76	0.539
C2F	23.63	73.89	0.333	25.51	90.62	0.277
EM	30.11	95.42	0.293	31.25	90.74	0.235
C2FSym	26.87	100.23	0.335	25.89	96.46	0.359
SM	28.81	68.42	0.326	32.66	75.38	0.240
ENC	23.60	51.32	0.385	29.29	57.28	0.303
RF [5]	22.26	69.26	0.548	26.92	79.43	0.435
Our approach	**4.49**	**10.96**	**0.907**	**5.95**	**14.18**	**0.858**

Figure 4: Evaluation on the FAUST dataset. For each technique, the table reports the average error on all pairs (AE, in centimeters), average error on the worst pair for this technique (worst AE, in centimeters), and the fraction of generated correspondences that are within 10cm of the ground truth (10cm-recall). Our technique reduces the average error by a factor of 3 over the best-performing prior approach on intra-subject pairs, and by a factor of 2.9 on inter-subject pairs. The worst AE is reduced by a factor of 4.7 on intra-subject pairs and by a factor of 4 on inter-subject pairs.

References

[1] Federica Bogo, Javier Romero, Matthew Loper, and Michael J. Black. FAUST: dataset and evaluation for 3D mesh registration. In *CVPR*, 2014.

[2] Alexander M. Bronstein, Michael M. Bronstein, and Ron Kimmel. Generalized multidimensional scaling: A framework for isometry-invariant partial surface matching. *Proceedings of the National Academy of Sciences*, 103(5), 2006.

[3] Vladimir G. Kim, Yaron Lipman, and Thomas A. Funkhouser. Blended intrinsic maps. *ACM Trans. Graph.*, 30(4), 2011.

[4] Vladimir Kolmogorov. Convergent tree-reweighted message passing for energy minimization. *PAMI*, 28(10), 2006.

[5] Emanuele Rodolà, Samuel Rota Bulò, Thomas Windheuser, Matthias Vestner, and Daniel Cremers. Dense non-rigid shape correspondence using random forests. In *CVPR*, 2014.