3D-CODED : 3D Correspondences by Deep Deformation

Thibault Groueix[1], Matthew Fisher[2], Vladimir G. Kim[2], Bryan C. Russell[2], Mathieu Aubry[1] [1]LIGM (UMR 8049), Ecole des Ponts, UPE, [2]Adobe Research
http://imagine.enpc.fr/~groueixt/3D-CODED/

Motivation

Task: 3D correspondences

Input: 3D pointclouds

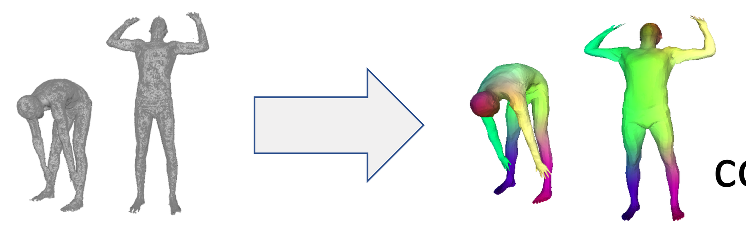

Output: 3D correspondences

Challenges:

- Low resolution sensors
- High sensor noise
- Expensive/impossible annotations

Contributions:

A simple framework

Previous work		**3D-CODED (ours)**
Manually designed template[5,2]	→	*any* neutral shape
Manual Parameterization	→	Learned Parameterization
Complex multiterm optimization	→	L2 loss + SGD

Strong unsupervised results

Supervised		**Unsupervised**
Per point correspondences	→	No annotation
L2 Distance	→	Chamfer Distance

Datasets: 230 000 synthetic human shapes

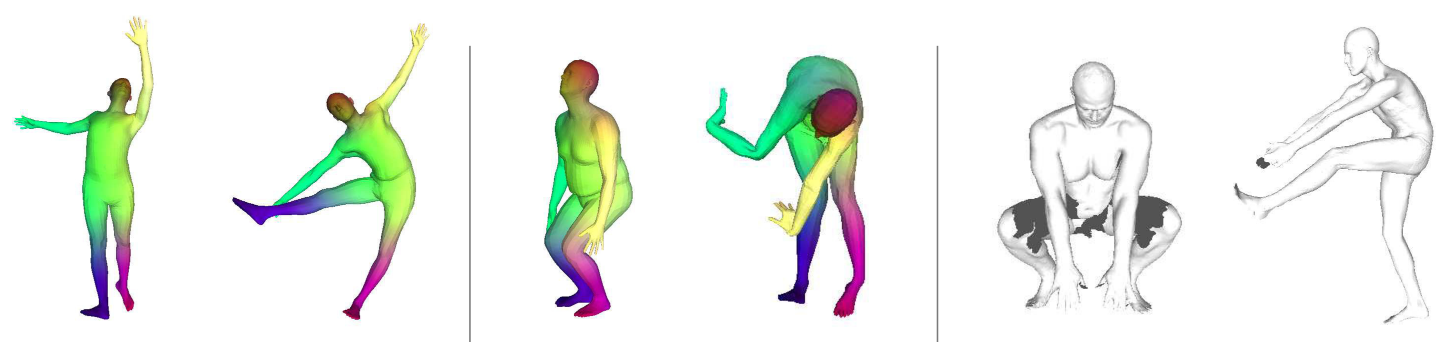

(a) SURREAL [1] (b) Bent shapes (c) FAUST [2]

synthetic training data (a, b), real testing data (c).

Key Ideas

Step 1: Learn 3D shape reconstruction by template deformation

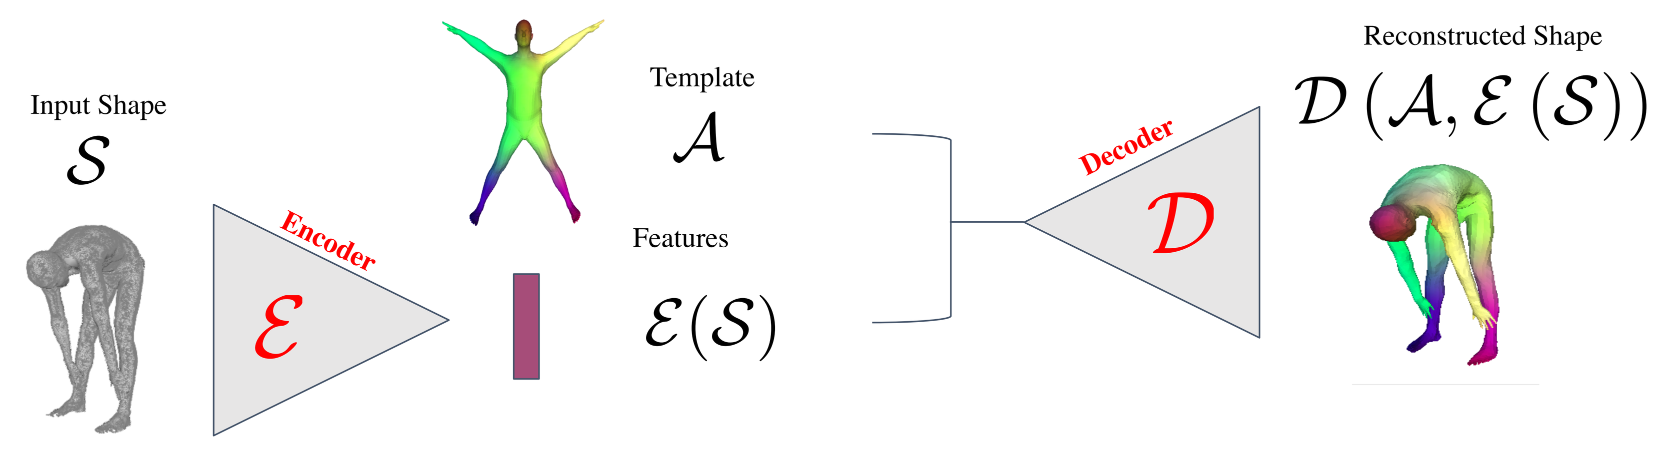

Supervised case: $\mathcal{L}^{\text{sup}}(\mathcal{E}, \mathcal{D}) = \sum_{j=1}^{\#points} |\mathcal{D}(\mathbf{p}_j; \mathcal{E}(\mathcal{S})) - \mathbf{q}_j|^2$

Unsupervised case: $\mathcal{L}^{\text{unsup}} = \mathcal{L}^{\text{CD}} + \lambda_{Lap}\mathcal{L}^{\text{Lap}} + \lambda_{edges}\mathcal{L}^{\text{edges}}$

$\mathcal{L}^{\text{CD}}$ **: Chamfer distance,** nearest neighbors based reconstruction loss

$\mathcal{L}^{\text{Lap}}$, $\mathcal{L}^{\text{edges}}$**: Laplacian loss, Edge loss** (regularization). Encourage isometry

Step 2: Optimizing shape reconstruction

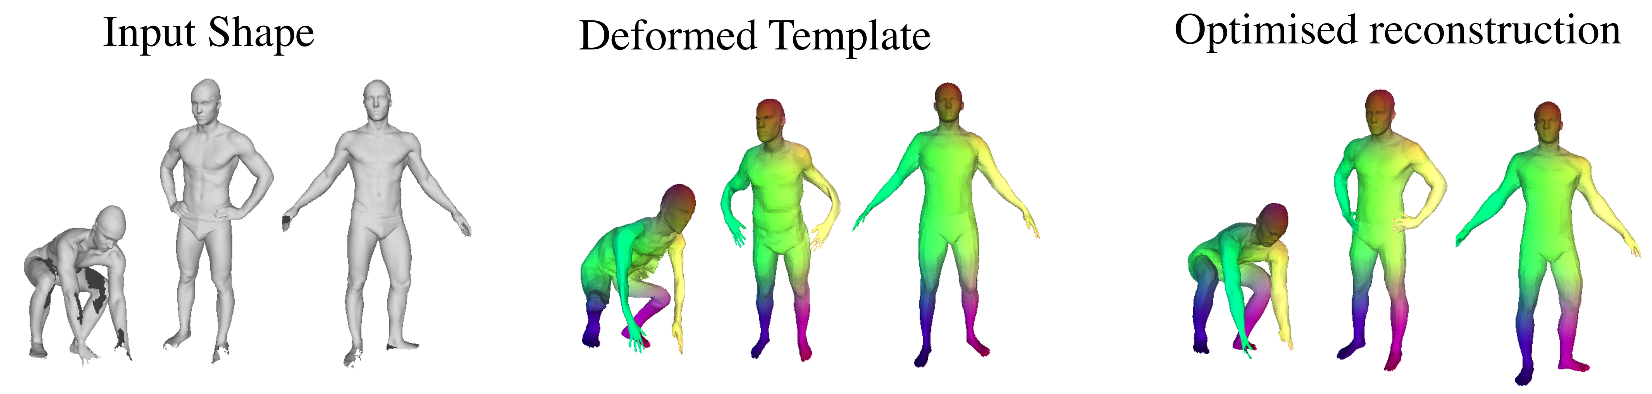

$$\mathbf{x}^* = \arg\min_{\mathbf{x}} \sum_{\mathbf{p}\in\mathcal{A}} \min_{\mathbf{q}\in\mathcal{S}} |\mathcal{D}(\mathbf{p};\mathbf{x}) - \mathbf{q}|^2 + \sum_{\mathbf{q}\in\mathcal{S}} \min_{\mathbf{p}\in\mathcal{A}} |\mathcal{D}(\mathbf{p};\mathbf{x}) - \mathbf{q}|^2 .$$

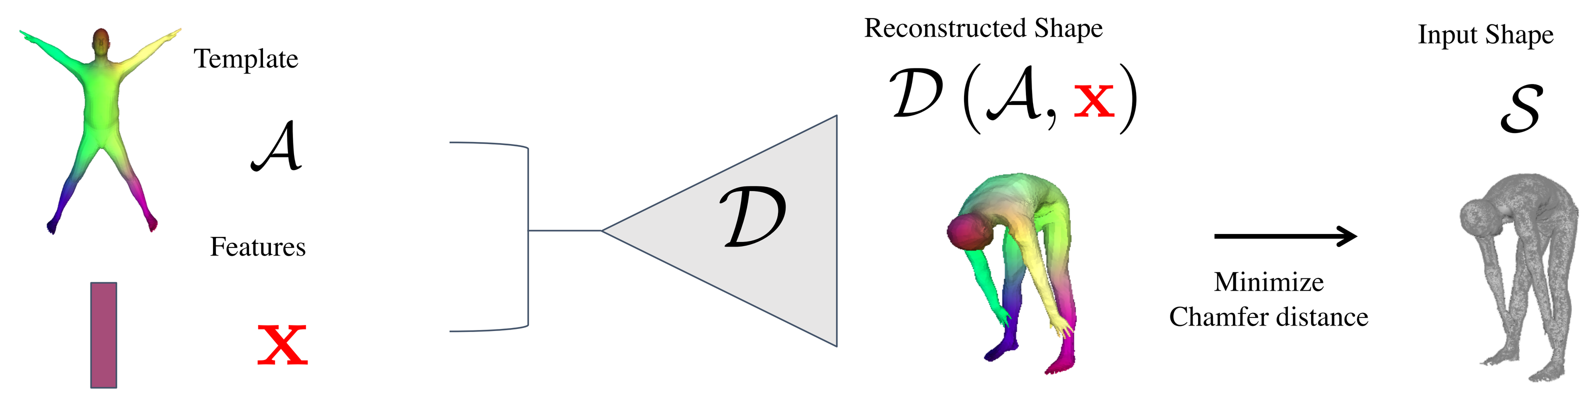

Step 3: Finding 3D shape correspondences

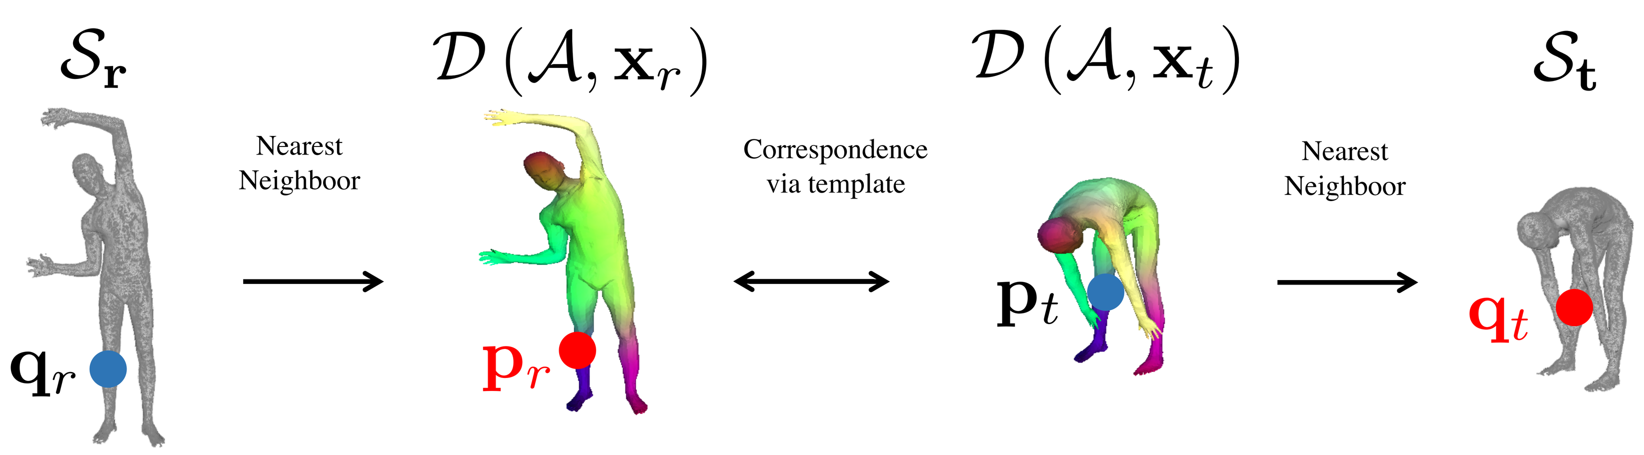

Results

State-of-the-art quantitative results

Method	Faust error
Convex-Opt [3]	8.304
FMNet [4]	4.826
SP [5]	3.126
Supervised	6.29
Supervised + Regression	3.255
Supervised + Regression + Regular Sampling	3.048
Supervised + Regression + Regular Sampling + High-Res template	**2.878**
Unsupervised + Regression + High-Res template	4.883

Faust Inter Challenge [2]. We outperform all other methods. Our unsupervised results are on par with other supervised methods. The reported error is the euclidean distance in cm.

Robustness to perturbations :

- noise, holes, sampling, topology, scaling

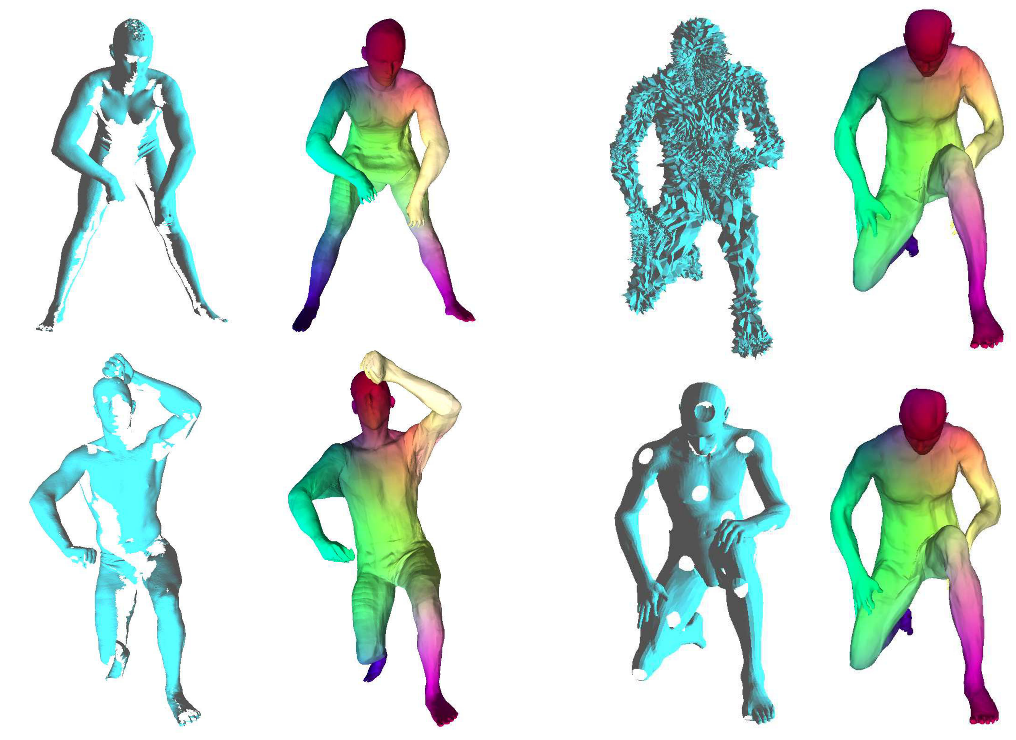

(a) SCAPE [6] (b) TOSCA [7]

Robustness to perturbations Left images show the input, right images the reconstruction with colors showing correspondences. Our method works with real incomplete scans **(a)**, and strong synthetic perturbations **(b)**.

Applicability beyond human

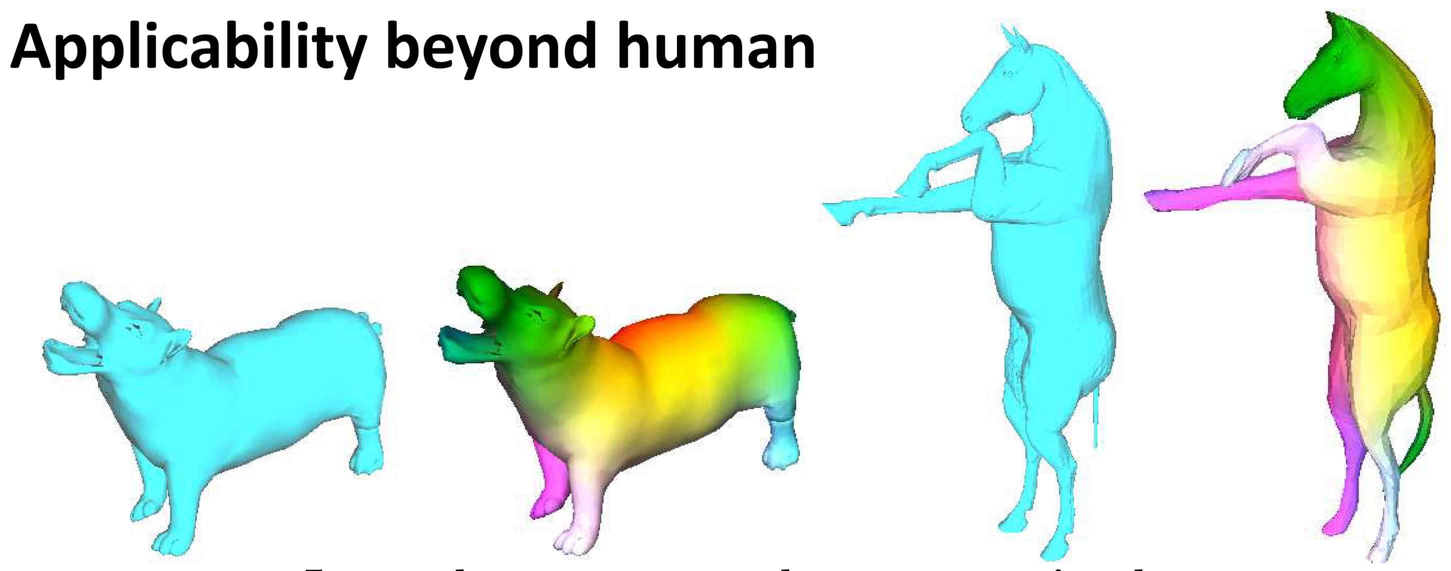

Inter-class correspondences on animals.

Code and results on the project webpage

[1] Learning from synthetic humans, Varol et al. CVPR (2017)
[2] FAUST: Dataset and evaluation for 3D mesh registration, Bogo et al. CVPR (2014)
[3] Robust nonrigid registration by convex optimization, Chen, Koltun, ICCV (2015)
[4] Deep functional maps: Structured prediction for dense shape correspondence, Litany et al. ICCV (2017)
[5] The stitched puppet: A graphical model of 3d human shape and pose, Zuffi et al. CVPR (2015)
[6] Scape: shape completion and animation of people, Anguelov et al. TOG (2005)
[7] Numerical geometry of non-rigid shapes, Bronstein et al. Springer (2008)