



3D-CODED: 3D Correspondences by Deep Deformation

Thibault Groueix¹(✉), Matthew Fisher², Vladimir G. Kim², Bryan C. Russell²,
and Mathieu Aubry¹

¹ LIGM (UMR 8049), École des Ponts, UPE, Champs-sur-Marne, France
thibault.groueix.2012@polytechnique.org

² Adobe Research, San Francisco, USA

<http://imagine.enpc.fr/~groueixt/3D-CODED/>

Abstract. We present a new deep learning approach for matching deformable shapes by introducing *Shape Deformation Networks* which jointly encode 3D shapes and correspondences. This is achieved by factoring the surface representation into (i) a template, that parameterizes the surface, and (ii) a learnt global feature vector that parameterizes the transformation of the template into the input surface. By predicting this feature for a new shape, we implicitly predict correspondences between this shape and the template. We show that these correspondences can be improved by an additional step which improves the shape feature by minimizing the Chamfer distance between the input and transformed template. We demonstrate that our simple approach improves on state-of-the-art results on the difficult FAUST-inter challenge, with an average correspondence error of 2.88 cm. We show, on the TOSCA dataset, that our method is robust to many types of perturbations, and generalizes to non-human shapes. This robustness allows it to perform well on real unclean, meshes from the SCAPE dataset.

Keywords: 3D deep learning · Computational geometry
Shape matching

1 Introduction

There is a growing demand for techniques that make use of the large amount of 3D content generated by modern sensor technology. An essential task is to establish reliable 3D shape correspondences between scans from raw sensor data or between scans and a template 3D shape. This process is challenging due to low sensor resolution and high sensor noise, especially for articulated shapes, such as humans or animals, that exhibit significant non-rigid deformations and shape variations (Fig. 1).

Electronic supplementary material The online version of this chapter (https://doi.org/10.1007/978-3-030-01216-8_15) contains supplementary material, which is available to authorized users.

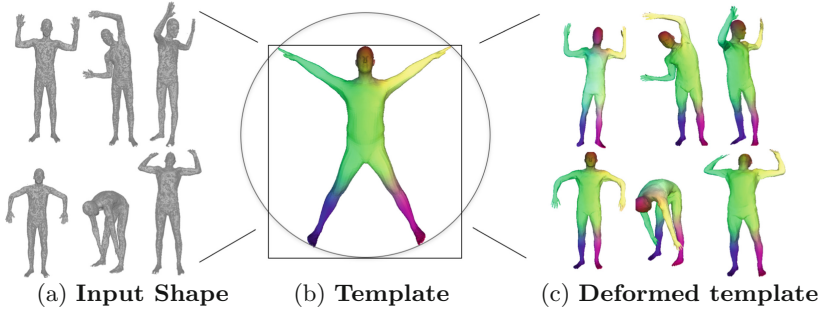


Fig. 1. Our approach predicts shape correspondences by learning a consistent mesh parameterization with a shared template. Colors show correspondences. (Color figure online)

Traditional approaches to estimating shape correspondences for articulated objects typically rely on intrinsic surface analysis either optimizing for an isometric map or leveraging intrinsic point descriptors [39]. To improve correspondence quality, these methods have been extended to take advantage of category-specific data priors [9]. Effective human-specific templates and registration techniques have been developed over the last decade [45], but these methods require significant effort and domain-specific knowledge to design the parametric deformable template, create an objective function that ensures alignment of salient regions and is not prone to being stuck in local minima, and develop an optimization strategy that effectively combines a global search for a good heuristic initialization and a local refinement procedure.

In this work, we propose *Shape Deformation Networks*, a comprehensive, all-in-one solution to template-driven shape matching. A Shape Deformation Network learns to deform a template shape to align with an input observed shape. Given two input shapes, we align the template to both inputs and obtain the final map between the inputs by reading off the correspondences from the template.

We train our Shape Deformation Network as part of an encoder-decoder architecture, which jointly learns an encoder network that takes a target shape as input and generates a global feature representation, and a decoder Shape Deformation Network that takes as input the global feature and deforms the template into the target shape. At test time, we improve our template-input shape alignment by optimizing locally the Chamfer distance between target and generated shape over the global feature representation which is passed in as input to the Shape Deformation Network. Critical to the success of our Shape Deformation Network is the ability to learn to deform a template shape to targets with varied appearances and articulation. We achieve this ability by training our network on a very large corpus of shapes.

In contrast to previous work [45], our method does not require a manually designed deformable template; the deformation parameters and degrees of freedom are implicitly learned by the encoder. Furthermore, while our network

can take advantage of known correspondences between the template and the example shapes, which are typically available when they have been generated using some parametric model [6, 41], we show it can also be trained without correspondence supervision. This ability allows the network to learn from a large collection of shapes lacking explicit correspondences.

We demonstrate that with sufficient training data this simple approach achieves state-of-the-art results and outperforms techniques that require complex multi-term objective functions instead of the simple reconstruction loss used by our method.

2 Related Work

Registration of non-rigid geometries with pose and shape variations is a long standing problem with extensive prior work. We first provide a brief overview of generic correspondence techniques. We then focus on category specific and template matching methods developed for human bodies, which are more closely related to our approach. Finally, we present an overview of deep learning approaches that have been developed for shape matching and more generally for working with 3D data.

Generic Shape Matching. To estimate correspondence between articulated objects, it is common to assume that their intrinsic structure (e.g., geodesic distances) remains relatively consistent across all poses [27]. Finding point-to-point correspondences that minimize metric distortion is a non-convex optimization problem, referred to as generalized multi-dimensional scaling [11]. This optimization is typically sensitive to an initial guess [10], and thus existing techniques rely on local feature point descriptors such as HKS [39] and WKS [5], and use hierarchical optimization strategies [14, 34]. Some relaxations of this problem have been proposed such as: formulating it as Markov random field and using linear programming relaxation [13], optimizing for soft correspondence [20, 37, 38], restricting correspondence space to conformal maps [21, 22], heat kernel maps [29], and aligning functional bases [30].

While these techniques are powerful generic tools, some common categories, such as humans, can benefit from a plethora of existing data [6] to leverage stronger class-specific priors.

Template-Based Shape Matching. A natural way to leverage class-specific knowledge is through the explicit use of a shape model. While such template-based techniques provide the best correspondence results they require a careful parameterization of the template, which took more than a decade of research to reach the current level of maturity [1–3, 24, 45]. For all of these techniques, fitting this representation to an input 3D shape requires also designing an objective function that is typically non-convex and involves multiple terms to guide the optimization to the right global minima. In contrast, our method only relies on

a single template 3D mesh and surface reconstruction loss. It leverages a neural network to learn how to parameterize the human body while optimizing for the best reconstruction.

Deep Learning for Shape Matching. Another way to leverage priors and training data is to learn better point-wise shape descriptors using human models with ground truth correspondence. Several neural network based methods have recently been developed to this end to analyze meshes [7, 26, 28, 33] or depth maps [42]. One can further improve these results by leveraging global context, for example, by estimating an inter-surface functional map [23]. These methods still rely on hand-crafted point-wise descriptors [40] as input and use neural networks to improve results. The resulting functional maps only align basis functions and additional optimization is required to extract consistent point-to-point correspondences [30]. One would also need to optimize for template deformation to use these matching techniques for surface reconstruction. In contrast our method does not rely on hand-crafted features (it only takes point coordinates as input) and implicitly learns a human body representation. It also directly outputs a template deformation.

Deep Learning for 3D Data. Following the success of deep learning approaches for image analysis, many techniques have been developed for processing 3D data, going beyond local descriptor learning to improve classification, segmentation, and reconstruction tasks. Existing networks operate on various shape representations, such as volumetric grids [17, 43], point clouds [16, 31, 32], geometry images [35, 36], seamlessly parameterized surfaces [25], by aligning a shape to a grid via distance-preserving maps [15], by folding a surface [44] or by predicting chart representations [18]. We build on these works in several ways. First, we process the point clouds representing the input shapes using an architecture similar to [31]. Second, similar to [35], we learn a surface representation. However, we do not explicitly encode correspondences in the output of a convolution network, but implicitly learn them by optimizing for parameters of the generation network as we optimize for reconstruction loss.

3 Method

Our goal is, given a reference shape \mathcal{S}_r and a target shape \mathcal{S}_t , to return a set of point correspondences \mathcal{C} between the shapes. We do so using two key ideas. First, we learn to predict a transformation between the shapes instead of directly learning the correspondences. This transformation, from 3D to 3D can indeed be represented by a neural network more easily than the association between variable and large number of points. The second idea is to learn transformations only from one template \mathcal{A} to any shape. Indeed, the large variety of possible poses of humans makes considering all pairs of possible poses intractable during training. We instead decouple the correspondence problem into finding two sets

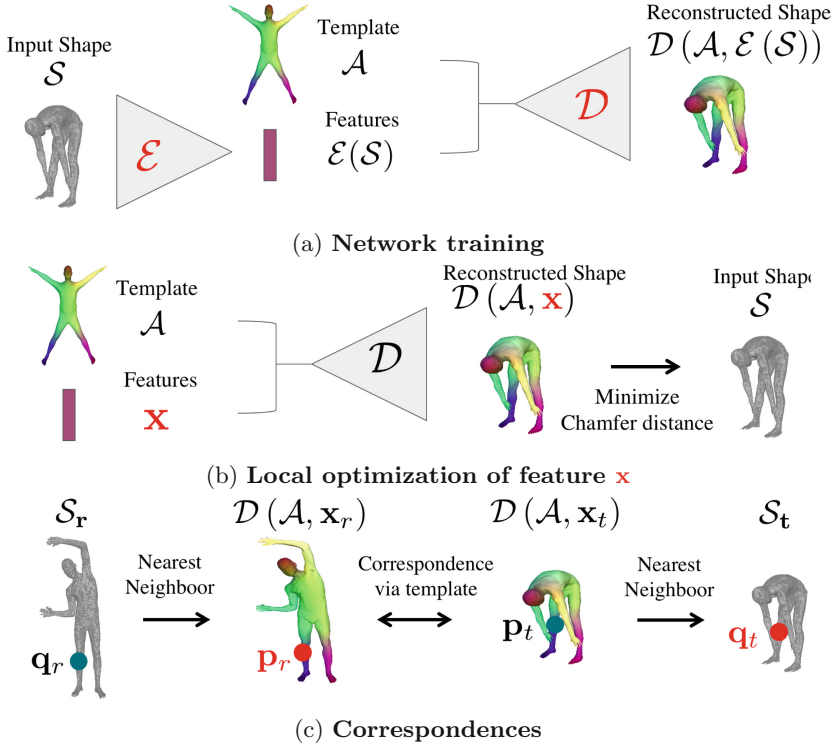


Fig. 2. Method overview. (a) A feed-forward pass in our autoencoder encodes input point cloud S to latent code $\mathcal{E}(S)$ and reconstruct S using $\mathcal{E}(S)$ to deform the template \mathcal{A} . (b) We refine the reconstruction $\mathcal{D}(\mathcal{A}, \mathcal{E}(S))$ by performing a regression step over the latent variable \mathbf{x} , minimizing the Chamfer distance between $\mathcal{D}(\mathcal{A}, \mathbf{x})$ and S . (c) Finally, given two point clouds S_r and S_t , to match a point \mathbf{q}_r on S_r to a point \mathbf{q}_t on S_t , we look for the nearest neighbor \mathbf{p}_r of \mathbf{q}_r in $\mathcal{D}(\mathcal{A}, \mathbf{x}_r)$, which is by design in correspondence with \mathbf{p}_t ; and look for the nearest neighbor \mathbf{q}_t of \mathbf{p}_t on S_t . Red indicates what is being optimised. (Color figure online)

of correspondences to a common template shape. We can then form our final correspondences between the input shapes via indexing through the template shape. An added benefit is during training we simply need to vary the pose for a single shape and use the known correspondences to the template shape as the supervisory signal.

Our approach has three main steps which are visualized Fig. 2. First, a feed-forward pass through our encoder network generates an initial global shape descriptor (Sect. 3.1). Second, we use gradient descent through our decoder Shape Deformation Network to refine this shape descriptor to improve the reconstruction quality (Sect. 3.2). We can then use the template to match points between any two input shapes (Sect. 3.3).

3.1 Learning 3D Shape Reconstruction by Template Deformation

To put an input shape \mathcal{S} in correspondence with a template \mathcal{A} , our first goal is to design a neural network that will take \mathcal{S} as input and predict transformation parameters. We do so by training an encoder-decoder architecture. The encoder \mathcal{E}_ϕ defined by its parameters ϕ takes as input 3D points, and is a simplified version of the network presented in [31]. It applies to each input 3D point coordinate a multi-layer perceptron with hidden feature size of 64, 128 and 1024, then maxpooling over the resulting features over all points followed by a linear layer, leading to feature of size 1024 $\mathcal{E}_\phi(\mathcal{S})$. This feature, together with the 3D coordinates of a point on the template $\mathbf{p} \in \mathcal{A}$, are taken as input to the decoder \mathcal{D}_θ with parameters θ , which is trained to predict the position \mathbf{q} of the corresponding point in the input shape. This decoder Shape Deformation Network is a multi-layer perceptron with hidden layers of size 1024, 512, 254 and 128, followed by a hyperbolic tangent. This architecture maps any points from the template domain to the reconstructed surface. By sampling the template more or less densely, we can generate an arbitrary number of output points by sequentially applying the decoder over sampled template points.

This encoder-decoder architecture is trained end-to-end. We assume that we are given as input a training set of N shapes $\{\mathcal{S}^{(i)}\}_{i=1}^N$ with each shape having a set of P vertices $\{\mathbf{q}_j\}_{j=1}^P$. We consider two training scenarios: one where the correspondences between the template and the training shapes are known (supervised case) and one where they are unknown (unsupervised case). Supervision is typically available if the training shapes are generated by deforming a parametrized template, but real object scans are typically obtained without correspondences.

Supervised Loss. In the supervised case, we assume that for each point \mathbf{q}_j on a training shape we know the correspondence $\mathbf{p}_j \leftrightarrow \mathbf{q}_j$ to a point $\mathbf{p}_j \in \mathcal{A}$ on the template \mathcal{A} . Given these training correspondences, we learn the encoder \mathcal{E}_ϕ and decoder \mathcal{D}_θ by simply optimizing the following reconstruction losses,

$$\mathcal{L}^{\text{sup}}(\theta, \phi) = \sum_{i=1}^N \sum_{j=1}^P |\mathcal{D}_\theta(\mathbf{p}_j; \mathcal{E}_\phi(\mathcal{S}^{(i)})) - \mathbf{q}_j^{(i)}|^2 \quad (1)$$

where the sums are over all P vertices of all N example shapes.

Unsupervised Loss. In the case where correspondences between the exemplar shapes and the template are not available, we also optimize the reconstructions, but also regularize the deformations toward isometries. For reconstruction, we use the Chamfer distance \mathcal{L}^{CD} between the inputs \mathcal{S}_i and reconstructed point clouds $\mathcal{D}_\theta(\mathcal{A}; \mathcal{E}_\phi(\mathcal{S}^{(i)}))$. For regularization, we use two different terms. The first term \mathcal{L}^{Lap} encourages the Laplacian operator defined on the template and the deformed template to be the same (which is the case for isometric deformations of the surface). The second term $\mathcal{L}^{\text{edges}}$ encourages the ratio between edges length

in the template and its deformed version to be close to 1. More details on these different losses are given in supplementary material. The final loss we optimize is:

$$\mathcal{L}^{\text{unsup}} = \mathcal{L}^{\text{CD}} + \lambda_{Lap} \mathcal{L}^{\text{Lap}} + \lambda_{edges} \mathcal{L}^{\text{edges}} \quad (2)$$

where λ_{Lap} and λ_{edges} control the influence of regularizations against the data term \mathcal{L}^{CD} . They are both set to $5 \cdot 10^{-3}$ in our experiments.

We optimize the loss using the Adam solver, with a learning rate of 10^{-3} for 25 epochs then 10^{-4} for 2 epochs, batches of 32 shapes, and 6890 points per shape.

One interesting aspect of our approach is that it learns jointly a parameterization of the input shapes via the decoder and to predict the parameters $\mathcal{E}_\phi(\mathcal{S})$ for this parameterization via the encoder. However, the predicted parameters $\mathcal{E}_\phi(\mathcal{S})$ for an input shape \mathcal{S} are not necessarily optimal, because of the limited power of the encoder. Optimizing these parameters turns out to be important for the final results, and is the focus of the second step of our pipeline.

3.2 Optimizing Shape Reconstruction

We now assume that we are given a shape \mathcal{S} as well as learned weights for the encoder \mathcal{E}_ϕ and decoder \mathcal{D}_θ networks. To find correspondences between the template shape and the input shape, we will use a nearest neighbor search to find correspondences between that input shape and its reconstruction. For this step to work, we need the reconstruction to be accurate. The reconstruction given by the parameters $\mathcal{E}_\phi(\mathcal{S})$ is only approximate and can be improved. Since we do not know correspondences between the input and the generated shape, we cannot minimize the loss given in Eq. (1), which requires correspondences. Instead, we minimize with respect to the global feature \mathbf{x} the Chamfer distance between the reconstructed shape and the input:

$$\mathcal{L}^{\text{CD}}(\mathbf{x}; \mathcal{S}) = \sum_{\mathbf{p} \in \mathcal{A}} \min_{\mathbf{q} \in \mathcal{S}} |\mathcal{D}_\theta(\mathbf{p}; \mathbf{x}) - \mathbf{q}|^2 + \sum_{\mathbf{q} \in \mathcal{S}} \min_{\mathbf{p} \in \mathcal{A}} |\mathcal{D}_\theta(\mathbf{p}; \mathbf{x}) - \mathbf{q}|^2. \quad (3)$$

Starting from the parameters predicted by our first step $\mathbf{x} = \mathcal{E}_\phi(\mathcal{S})$, we optimize this loss using the Adam solver for 3,000 iterations with learning rate $5 * 10^{-4}$. Note that the good initialization given by our first step is key since Eq. (3) corresponds to a highly non-convex problem, as shown in Fig. 6.

3.3 Finding 3D Shape Correspondences

To recover correspondences between two 3D shapes \mathcal{S}_r and \mathcal{S}_t , we first compute the parameters to deform the template to these shapes, \mathbf{x}_r and \mathbf{x}_t , using the two steps outlined in Sects. 3.1 and 3.2. Next, given a 3D point \mathbf{q}_r on the reference shape \mathcal{S}_r , we first find the point \mathbf{p} on the template \mathcal{A} such that its transformation with parameters \mathbf{x}_r , $\mathcal{D}_\theta(\mathbf{p}; \mathbf{x}_r)$ is closest to \mathbf{q}_r . Finally we find the 3D point \mathbf{q}_t on the target shape \mathcal{S}_t that is the closest to the transformation of \mathbf{p} with parameters \mathbf{x}_t , $\mathcal{D}_\theta(\mathbf{p}; \mathbf{x}_t)$. Our algorithm is summarized in Algorithm 1 and illustrated in Fig. 2.

Algorithm 1. Algorithm for finding 3D shape correspondences

Input : Reference shape \mathcal{S}_r and target shape \mathcal{S}_t
Output: Set of 3D point correspondences \mathcal{C}

- 1 #Regression steps over latent code to find best reconstruction of \mathcal{S}_r and \mathcal{S}_t
- 2 $\mathbf{x}_r \leftarrow \arg \min_{\mathbf{x}} \mathcal{L}^{\text{CD}}(\mathbf{x}; \mathcal{S}_r)$ #detailed in equation (3)
- 3 $\mathbf{x}_t \leftarrow \arg \min_{\mathbf{x}} \mathcal{L}^{\text{CD}}(\mathbf{x}; \mathcal{S}_t)$ #detailed in equation (3)
- 4 $\mathcal{C} \leftarrow \emptyset$
- 5 # Matching of $\mathbf{q}_r \in \mathcal{S}_r$ to $\mathbf{q}_t \in \mathcal{S}_t$
- 6 **foreach** $\mathbf{q}_r \in \mathcal{S}_r$ **do**
- 7 $\mathbf{p} \leftarrow \arg \min_{\mathbf{p}' \in \mathcal{A}} |\mathcal{D}_\theta(\mathbf{p}'; \mathbf{x}_r) - \mathbf{q}_r|^2$
- 8 $\mathbf{q}_t \leftarrow \arg \min_{\mathbf{q}' \in \mathcal{S}_t} |\mathcal{D}_\theta(\mathbf{p}; \mathbf{x}_t) - \mathbf{q}'|^2$
- 9 $\mathcal{C} \leftarrow \mathcal{C} \cup \{(\mathbf{q}_r, \mathbf{q}_t)\}$
- 10 **end**
- 11 **return** \mathcal{C}

4 Results

4.1 Datasets

Synthetic Training Data. To train our algorithm, we require a large set of shapes. We thus rely on synthetic data for training our model.

For human shapes, we use SMPL [6], a state-of-the-art generative model for synthetic humans. To obtain realistic human body shape and poses from the SMPL model, we sampled 2.10^5 parameters estimated in the SURREAL dataset [41]. One limitation of the SURREAL dataset is it does not include any humans bent over. Without adapted training data, our algorithm generalized poorly to these poses. To overcome this limitation, we generated an extension of the dataset. We first manually estimated 7 key-joint parameters (among 23 joints in the SMPL skeletons) to generate bent humans. We then sampled randomly the 7 parameters around these values, and used parameters from the SURREAL dataset for the other pose and body shape parameters. Note that not all meshes generated with this strategy are realistic as shown in Fig. 3. They however allow us to better cover the space of possible poses, and we added $3 \cdot 10^4$ shapes generated with this method to our dataset. Our final dataset thus has $2.3 \cdot 10^5$ human meshes with a large variety of realistic poses and body shapes.

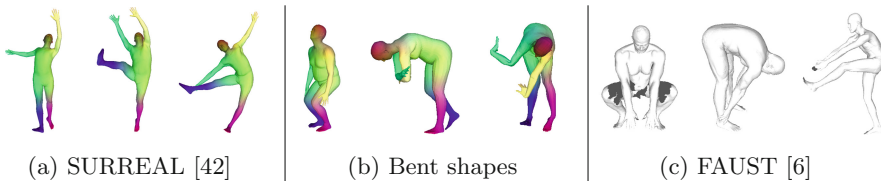


Fig. 3. Examples of the different datasets used in the paper.

For animal shapes, we use the SMAL [47] model, which provides the equivalent of SMPL for several animals. Recent papers estimate model parameters from images, but no large-scale parameter set is yet available. For training we thus generated models from SMAL with random parameters (drawn from a Gaussian distribution of *ad-hoc* variance 0.2). This approach works for the 5 categories available in SMAL. In SMALR [46], Zuffi et al. showed that the SMAL model could be generalized to other animals using only an image dataset as input, demonstrating it on 17 additional categories. Note that since the templates for two animals are in correspondences, our method can be used to get inter-category correspondences for animals. We qualitatively demonstrate this on hippopotamus/horses in the appendix [19].

Testing Data. We evaluate our algorithm on the FAUST [6], TOSCA [12] and SCAPE [4] datasets.

The FAUST dataset consists of 100 training and 200 testing scans of approximately 170,000 vertices. They may include noise and have holes, typically missing part of the feet. In this paper, we never used the training set, except for a single baseline experiment, and we focus on the test set. Two challenges are available, focusing on intra- and inter-subject correspondences. The error is the average Euclidean distance between the estimated projection and the ground-truth projection. We evaluated our method through the online server and are the best public results on the ‘inter’ challenge at the time of submission¹.

The SCAPE [4] dataset has two sets of 71 meshes: the first set consists of real scans with holes and occlusions and the second set are registered meshes aligned to the first set. The poses are different from both our training dataset and FAUST.

TOSCA is a dataset produced by deforming 3 template meshes (human, dog, and horse). Each mesh is deformed into multiple poses, and might have various additional perturbations such as random holes in the surface, local and global scale variations, noise in vertex positions, varying sampling density, and changes in topology.

Shape Normalization. To be processed and reconstructed by our network, the training and testing shapes must be normalized in a similar way. Since the vertical direction is usually known, we used synthetic shapes with approximately the same vertical axis. We also kept a fixed orientation around this vertical axis, and at test time selected the one out of 50 different orientations which leads to the smaller reconstruction error in term of Chamfer distance. Finally, we centered all meshes according to the center of their bounding box and, for the training data only, added a random translation in each direction sampled uniformly between -3 cm and 3 cm to increase robustness.

¹ <http://faust.is.tue.mpg.de/challenge/Inter-subject.challenge>.

4.2 Experiments

In this part, we analyze the key components of our pipeline. More results are available in the appendix [19].

Results on FAUST. The method presented above leads to the best results to date on the FAUST-inter dataset: 2.878 cm: **an improvement of 8% over state of the art**, 3.12 cm for [45] and 4.82 cm for [23]. Although it cannot take advantage of the fact that two meshes represent the same person, our method is also the second best performing (average error of 1.99 cm) on FAUST-intra challenge.

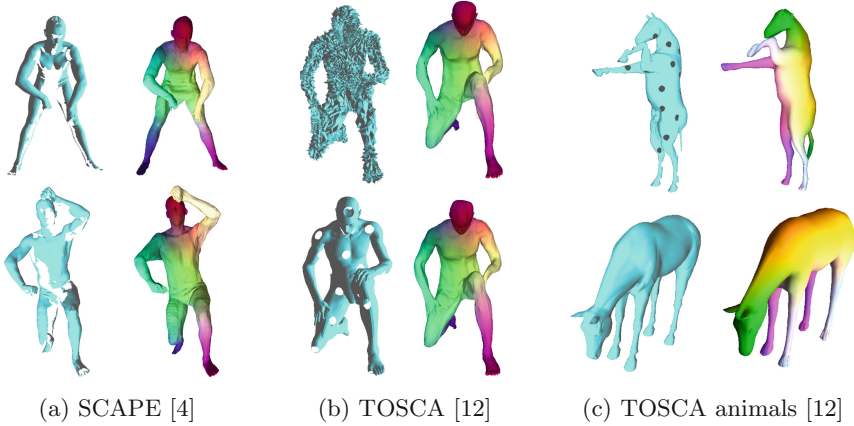


Fig. 4. Other datasets. Left images show the input, right images the reconstruction with colors showing correspondences. Our method works with real incomplete scans (a), strong synthetic perturbations (b), and on non-human shapes (c). (Color figure online)

Results on SCAPE: Real and Partial Data. The SCAPE dataset provides meshes aligned to real scans and includes poses different from our training dataset. When applying a network trained directly on our SMPL data, we obtain satisfying performance, namely 3.14 cm average Euclidean error. Quantitative comparison of correspondence quality in terms of geodesic error are given in Fig. 5. We outperform all methods except for Deep Functional Maps [23]. SCAPE also allows evaluation on real partial scans. Quantitatively, the error on these partial meshes is 4.04 cm, similar to the performance on the full meshes. Qualitative results are shown in Fig. 4a.

Results on TOSCA: Robustness to Perturbations. The TOSCA dataset provides several versions of the same synthetic mesh with different perturbations. We found that our method, still trained only on SMPL or SMAL data, is

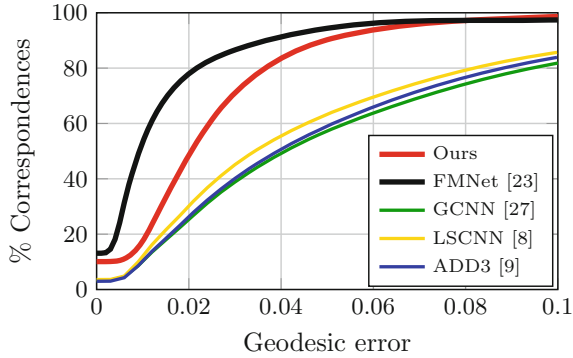


Fig. 5. Comparison with learning-based shape matching approaches on the SCAPE dataset. Our method is trained on synthetic data, FMNet was trained on FAUST data, and all other methods on SCAPE. We outperform all methods except FMNet even though our method was trained on a different dataset.

robust to all perturbations (isometry, noise, shotnoise, holes, micro-holes, topology changes, and sampling), except scale, which can be trivially fixed by normalizing all meshes to have consistent surface area. Examples of representative qualitative results are shown Fig. 4 and quantitative results are reported in appendix [19].

Table 1. Importance of the reconstruction optimization step. Optimizing the latent feature is key to our results. Regular point sampling for training and high resolution for the nearest neighbor step provide an additional boost.

Method	Faust error (cm)
Without regression	6.29
With regression	3.255
With regression + regular sampling	3.048
With regression + regular sampling + high-res template	2.878

Reconstruction Optimization. Because the nearest neighbors used in the matching step are sensitive to small errors in alignment, the second step of our pipeline which finds the optimal features for reconstruction, is crucial to obtain high quality results. This optimization however converges to a good optimum only if it is initialized with a reasonable reconstruction, as visualized in Fig. 6. Since we optimize using Chamfer distance, and not correspondences, we also rely on the fact that the network was trained to generate humans in correspondence and we expect the optimized shape to still be meaningful.

Table 1 reports the associated quantitative results on FAUST-inter. We can see that: (i) optimizing the latent feature to minimize the Chamfer distance

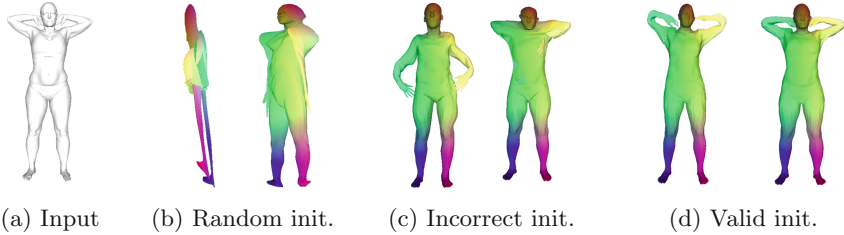


Fig. 6. Reconstruction optimization. The quality of the initialization (i.e. the first step of our algorithm) is crucial for the deformation optimization. For a given target shape (a) and for different initializations (left of (b), (c) and (d)) the figure shows the results of the optimization. If the initialization is random (b) or incorrect (c), the optimization converges to bad local minima. With a reasonable initialization (d) it converges to a shape very close to the target ((d), right).

between input and output provides a strong boost; (ii) using a better (more uniform) sampling of the shapes when training our network provided a better initialization; (iii) using a high resolution sampling of the template ($\sim 200k$ vertices) for the nearest-neighbor step provide an additional small boost in performance.

Table 2. FAUST-inter results when training on different datasets. Adding synthetic data reduce the error by a factor of 3, showing its importance. The difference in performance between the basic synthetic dataset and its augmented version is mostly due to failure on specific poses, as in Fig. 3.

Training data	Faust error (cm)
FAUST training set	18.22
Non-augmented synthetic dataset 2×10^5 shapes	5.63
Augmented synthetic data, 10^3 shapes	5.76
Augmented synthetic data, 10^4 shapes	4.70
Augmented synthetic data, 2.3×10^5 shapes	3.26

Necessary Amount of Training Data. Training on a large and representative dataset is also crucial for our method. To analyze the effect of training data, we ran our method without re-sampling FAUST points regularly and with a low resolution template for different training sets: FAUST training set, 2×10^5 SURREAL shapes, and 2.3×10^5 , 10^4 and 10^3 shapes from our augmented dataset. The quantitative results are reported Table 2 and qualitative results can be seen in Fig. 7. The FAUST training set only include 10 different poses and is too small to train our network to generalize. Training on many synthetic shapes from the SURREAL dataset [41] helps overcome this generalization problem. However, if the synthetic dataset does not include any pose close to test poses (such as




Fig. 7. Importance of the training data. For a given target shape (a) reconstructed shapes when the network is trained on FAUST training set (b) and on our augmented synthetic training set (c), before (left) and after (right) the optimization step.

bent-over humans), the method will fail on these poses (4 test pairs of shapes out of 40). Augmenting the dataset as described in Sect. 4.1 overcomes this limitation. As expected the performance decreases with the number of training shapes, respectively to 5.76 cm and 4.70 cm average error on FAUST-inter.

Table 3. Results with and without supervised correspondences. Adding regularization helps the network find a better local minimum in terms of correspondences.

Loss	Faust error (cm)
Chamfer distance, Eq. 3 (unsupervised)	8.727
Chamfer distance + Regularization, Eq. 2 (unsupervised)	4.835
Correspondences, Eq. 1 (supervised)	2.878

Unsupervised Correspondences. We investigate whether our method could be trained without correspondence supervision. We started by simply using the reconstruction loss described in Eq. (3). One could indeed expect that an optimal way to deform the template into training shapes would respect correspondences. However, we found that the resulting network did not respect correspondences between the template and the input shape, as visualized Fig. 8. However, these results improve with adequate regularization such as the one presented in Eq. (2), encouraging regularity of the mapping between the template and the reconstruction. We trained such a network with the same training data as in the supervised case but **without any correspondence supervision** and obtained a 4.88 cm of error on the FAUST-inter data, i.e. similar to Deep Functional Map [23] which had an error of 4.83 cm. This demonstrates that our method can be efficient even without correspondence supervision. Further details on regularization losses are given in the appendix [19] Table 3.

Rotation Invariance. We handled rotation invariance by rotating the shape and selecting the orientation for which the reconstruction is optimal. As an alternative, we tried to learn a network directly invariant to rotations around

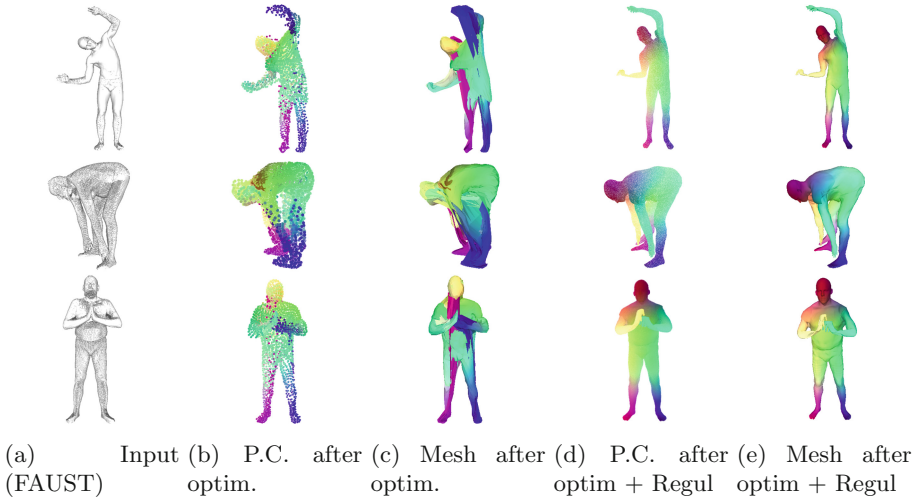


Fig. 8. Unsupervised correspondences. We visualize for different inputs (a), the point clouds (P.C.) predicted by our approach (b, d) and the corresponding meshes (c, e). Note that without regularization, because of the strong distortion, the meshes appear to barely match to the input, while the point clouds are reasonable. On the other hand surface regularization creates reasonable meshes.

the vertical axis. It turned out the performances were slightly worse on FAUST-inter (3.10 cm), but still better than the state of the art. We believe this is due to the limited capacity of the network and should be tried with a larger network. However, interestingly, this rotation invariant network seems to have increased robustness and provided slightly better results on SCAPE.

5 Conclusion

We have demonstrated an encoder-decoder deep network architecture that can generate human shape correspondences competitive with state-of-the-art approaches and that uses only simple reconstruction and correspondence losses. Our key insight is to factor the problem into an encoder network that produces a global shape descriptor, and a decoder Shape Deformation Network that uses this global descriptor to map points on a template back to the original geometry. A straightforward regression step uses gradient descent through the Shape Deformation Network to significantly improve the final correspondence quality.

Acknowledgments. This work was partly supported by ANR project EnHerit ANR-17-CE23-0008, Labex Bézout, and gifts from Adobe to École des Ponts. We thank Gül Varol, Angjoo Kanazawa, and Renaud Marlet for fruitful discussions.

References

1. Allen, B., Curless, B., Popovic, Z.: Articulated body deformation from range scan data. In: SIGGRAPH (2002)
2. Allen, B., Curless, B., Popovic, Z.: The space of human body shapes: reconstruction and parameterization from range scans. In: SIGGRAPH (2003)
3. Allen, B., Curless, B., Popovic, Z.: Learning a correlated model of identity and pose-dependent body shape variation for real-time synthesis. In: Symposium on Computer Animation (2006)
4. Anguelov, D., Srinivasan, P., Koller, D., Thrun, S., Rodgers, J., Davis, J.: SCAPE: shape completion and animation of people. *ACM Trans. Graph. (TOG)* **24**(3), 408–416 (2005)
5. Aubry, M., Schlickewei, U., Cremers, D.: The wave kernel signature: a quantum mechanical approach to shape analysis. In: IEEE International Conference on Computer Vision (ICCV) - Workshop on Dynamic Shape Capture and Analysis (4DMOD) (2011)
6. Bogo, F., Romero, J., Loper, M., Black, M.J.: FAUST: dataset and evaluation for 3D mesh registration. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), June 2014
7. Boscaini, D., Masci, J., Rodola, E., Bronstein, M.M.: Learning shape correspondence with anisotropic convolutional neural networks. In: NIPS (2016)
8. Boscaini, D., Masci, J., Melzi, S., Bronstein, M.M., Castellani, U., Vanderghenst, P.: Learning class-specific descriptors for deformable shapes using localized spectral convolutional networks. *Comput. Graph. Forum* **34**(5), 13–23 (2015)
9. Boscaini, D., Masci, J., Rodola, E., Bronstein, M.M., Cremers, D.: Anisotropic diffusion descriptors. *Comput. Graph. Forum* **35**(2), 431–441 (2016)
10. Bronstein, A.M., Bronstein, M.M., Kimmel, R.: Efficient computation of isometry-invariant distances between surfaces. *SIAM J. Sci. Comput.* **28**(5), 1812–1836 (2006)
11. Bronstein, A.M., Bronstein, M.M., Kimmel, R.: Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. In: Proceedings of the National Academy of Sciences (PNAS) (2006)
12. Bronstein, A.M., Bronstein, M.M., Kimmel, R.: *Numerical Geometry of Non-Rigid Shapes*. MCS. Springer Science & Business Media, New York (2009). <https://doi.org/10.1007/978-0-387-73301-2>
13. Chen, Q., Koltun, V.: Robust nonrigid registration by convex optimization. In: International Conference on Computer Vision (ICCV) (2015)
14. Raviv, D., Dubrovina, A., Kimmel, R.: Hierarchical framework for shape correspondence. *Numer. Math. Theory Methods Appl.* (2013)
15. Ezuz, D., Solomon, J., Kim, V.G., Ben-Chen, M.: GWCNN: a metric alignment layer for deep shape analysis. In: SGP (2017)
16. Fan, H., Su, H., Guibas, L.: A point set generation network for 3D object reconstruction from a single image. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2017)
17. Girdhar, R., Fouhey, D.F., Rodriguez, M., Gupta, A.: Learning a predictable and generative vector representation for objects. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016. LNCS, vol. 9910, pp. 484–499. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46466-4_29
18. Groueix, T., Fisher, M., Kim, V.G., Russell, B., Aubry, M.: AtlasNet: a Papier-Mâché approach to learning 3D surface generation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018)

19. Groueix, T., Fisher, M., Kim, V.G., Russell, B., Aubry, M.: Supplementary material (appendix) for the paper (2018). <http://imagine.enpc.fr/~groueix/3D-CODED/index.html>
20. Kim, V.G., Li, W., Mitra, N.J., DiVerdi, S., Funkhouser, T.: Exploring collections of 3D models using fuzzy correspondences. *Trans. Graph.* **31**(4), 54:1–54:11 (2012). (Proc. of SIGGRAPH)
21. Kim, V.G., Lipman, Y., Funkhouser, T.: Blended intrinsic maps. *Trans. Graph.* **30**(4) (2011). (Proc. of SIGGRAPH)
22. Lipman, Y., Funkhouser, T.: Mobius voting for surface correspondence. *ACM Trans. Graph.* **28**(3) (2009). (Proc. SIGGRAPH)
23. Litany, O., Remez, T., Rodola, E., Bronstein, A.M., Bronstein, M.M.: Deep functional maps: structured prediction for dense shape correspondence. In: *ICCV* (2017)
24. Loper, M., Mahmood, N., Romero, J., Pons-Moll, G., Black, M.J.: SMPL: a skinned multi-person linear model. In: *SIGGRAPH Asia* (2015)
25. Maron, H., et al.: Convolutional neural networks on surfaces via seamless toric covers. In: *SIGGRAPH* (2017)
26. Masci, J., Boscaini, D., Bronstein, M.M., Vandergheynst, P.: Geodesic convolutional neural networks on riemannian manifolds. In: *Proceedings of the IEEE International Conference on Computer Vision (ICCV) Workshops*, pp. 37–45 (2015)
27. Mémoli, F., Sapiro, S.: A theoretical and computational framework for isometry invariant recognition of point cloud data. *Found. Comput. Math.* **5**(3), 313–347 (2005)
28. Monti, F., Boscaini, D., Masci, J., Rodola, E., Svoboda, J., Bronstein, M.M.: Geometric deep learning on graphs and manifolds using mixture model CNNs. In: *CVPR* (2017)
29. Ovsjanikov, M., Mérigot, Q., Mémoli, F., Guibas, L.: One point isometric matching with the heat kernel. *Comput. Graph. Forum* **29**, 1555–1564 (2010). (Proc. of SGP)
30. Ovsjanikov, M., Ben-Chen, M., Solomon, J., Butscher, A., Guibas, L.: Functional maps: a flexible representation of maps between shapes. *ACM Trans. Graph.* (2012)
31. Qi, C.R., Su, H., Mo, K., Guibas, L.J.: PointNet: deep learning on point sets for 3D classification and segmentation. In: *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2017)
32. Qi, C.R., Yi, L., Su, H., Guibas, L.J.: PointNet++: deep hierarchical feature learning on point sets in a metric space. In: *Advances in Neural Information Processing Systems (NIPS)* (2017)
33. Rodola, E., Rota Bulo, S., Windheuser, T., Vestner, M., Cremers, D.: Dense non-rigid shape correspondence using random forests. In: *CVPR* (2014)
34. Sahillioglu, Y., Yemez, Y.: Coarse-to-fine combinatorial matching for dense isometric shape correspondence. *Comput. Graph. Forum* **30**, 1461–1470 (2011)
35. Sinha, A., Unmesh, A., Huang, Q., Ramani, K.: SurfNet: generating 3D shape surfaces using deep residual networks. In: *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2017)
36. Sinha, A., Bai, J., Ramani, K.: Deep learning 3D shape surfaces using geometry images. In: *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2016)
37. Solomon, J., Nguyen, A., Butscher, A., Ben-Chen, M., Guibas, L.: Soft maps between surfaces. In: *SGP* (2012)
38. Solomon, J., Peyre, G., Kim, V.G., Sra, S.: Entropic metric alignment for correspondence problems. *Trans. Graph.* **35**, 1–13 (2016). (Proc. of SIGGRAPH)

39. Sun, J., Ovsjanikov, M., Guibas, L.: A concise and provably informative multi-scale signature-based on heat diffusion. *Comput. Graph. Forum* **28**, 1383–1392 (2009). (Proc. of SGP)
40. Tombari, F., Salti, S., Di Stefano, L.: Unique signatures of histograms for local surface description. In: Daniilidis, K., Maragos, P., Paragios, N. (eds.) *ECCV 2010*. LNCS, vol. 6313, pp. 356–369. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-15558-1_26
41. Varol, G., et al.: Learning from synthetic humans. In: *CVPR (2017)*
42. Wei, L., Huang, Q., Ceylan, D., Vouga, E., Li, H.: Dense human body correspondences using convolutional networks. In: *Computer Vision and Pattern Recognition (CVPR) (2016)*
43. Wu, Z., et al.: 3D ShapeNets: a deep representation for volumetric shapes. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1912–1920 (2015)
44. Yang, Y., Feng, C., Shen, Y., Tian, D.: FoldingNet: point cloud auto-encoder via deep grid deformation. In: *CVPR (2018)*
45. Zuffi, S., Black, M.J.: The stitched puppet: a graphical model of 3D human shape and pose. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2015)*
46. Zuffi, S., Kanazawa, A., Black, M.J.: Lions and tigers and bears: capturing non-rigid, 3D, articulated shape from images. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018)*
47. Zuffi, S., Kanazawa, A., Jacobs, D., Black, M.J.: 3D menagerie: modeling the 3D shape and pose of animals. In: *CVPR (2017)*