Published in final edited form as:

IEEE Trans Pattern Anal Mach Intell. 2018 December ; 40(12): 3059–3066. doi:10.1109/TPAMI.2017.2772922.

A Simple, Fast and Highly-Accurate Algorithm to Recover 3D Shape from 2D Landmarks on a Single Image

Ruiqi Zhao, Yan Wang, and Aleix M. Martinez

Department of Electrical and Computer Engineering, The Ohio State University, Columbus, OH, 43210. <http://cbcs1.ece.ohio-state.edu>

Abstract

Three-dimensional shape reconstruction of 2D landmark points on a single image is a hallmark of human vision, but is a task that has been proven difficult for computer vision algorithms. We define a feed-forward deep neural network algorithm that can reconstruct 3D shapes from 2D landmark points almost perfectly (i.e., with extremely small reconstruction errors), even when these 2D landmarks are from a single image. Our experimental results show an improvement of up to two-fold over state-of-the-art computer vision algorithms; 3D shape reconstruction error (measured as the Procrustes distance between the reconstructed shape and the ground-truth) of human faces is $< .004$, cars is $.0022$, human bodies is $.022$, and highly-deformable flags is $.0004$. Our algorithm was also a top performer at the 2016 3D Face Alignment in the Wild Challenge competition (done in conjunction with the European Conference on Computer Vision, ECCV) that required the reconstruction of 3D face shape from a single image. The derived algorithm can be trained in a couple hours and testing runs at more than 1,000 frames/s on an i7 desktop. We also present an innovative data augmentation approach that allows us to train the system efficiently with small number of samples. And the system is robust to noise (e.g., imprecise landmark points) and missing data (e.g., occluded or undetected landmark points).

Keywords

3D modeling and reconstruction; fine-grained reconstruction; 3D shape from a single 2D image; deep learning

1 Introduction

Humans can readily and accurately estimate the 3D shape of an object from a set of 2D landmark points on a single image [1]. Many computer vision approaches have been developed over the years in an attempt to replicate this outstanding human ability. As an example, the approaches in [2], [3], [4], [5], [6] provide 3D shape estimates of a set of 2D landmark points on a single image. Other works, such as structure from motion (SfM) and shape from shading, require a sequence of images [7], [8], [9]. In general, the results of SfM are not as accurate as the ones reported in the present paper. Although less precise, an advantage of SfM is that they do not require learning from labelled data and can work on most object classes.

The above mentioned algorithms learn a set of 3D shape bases, with the assumption that the deformation in a test image can be represented as a linear combination of these bases. This linear assumption limits the applicability of these approaches on highly deformable or articulated shapes, a point that is only exacerbated when the number of 2D landmark points is small (i.e., non-dense fiducials). To resolve these problems, some previous works, e.g., [3], [4], are specifically designed to recover the 3D shape of an object by introducing domain-specific priors into the 3D estimate. However, the addition of these priors typically limit the applicability of the resulting algorithms. Another solution is to identify the set of piecewise segments that belong to the same surface and then solve for each section separately [10], [11]. However, these algorithms require multiple images or video sequences to identify consistency and smoothness in the movement.

Estimating the 3D geometry of an object from a single view is an ill-posed problem. Nevertheless, with the available 3D ground-truth of a number of 2D sample images with corresponding 2D landmarks, we can learn the mapping from 2D to 3D landmark points, i.e., we can define a mapping function $f(\cdot)$ that given a set of 2D landmark points x outputs their corresponding 3D coordinates y , $y = f(x)$. There are three main challenges in this task: (1) how to define a general algorithm that works in a wide range of rigid *and* non-rigid objects; (2) how to define an algorithm that yields good results whether we have a small or a large number of samples; and (3) how to make this process run in real-time (> 60 frames/s).

In order to deal with the aforementioned challenges, we propose a deep-network framework to estimate the function $f(\cdot)$. The proposed model is illustrated in Figure 1. Unlike most existing methods, our work derives a deep neural network that uses a set of hierarchical nonlinear transformations to recover depth information of a number of known 2D landmark points on a single image. The 3D shape is estimated by combining the input and output of the neural network and the shape is recovered up to scale.

The major contributions of the herein derived algorithm can be summarized as follows.

1. Our algorithm is extremely general and can be applied to recover the 3D shape of basically any rigid or non-rigid object for which a number of training samples (i.e., 2D and their corresponding 3D landmark points) is available. As examples, we provide results on human faces, cars, human bodies and flags.
2. Our algorithm is not limited by the number of training samples. When the number of training data is very large, we employ mini-batch training to parallelize gradient-based optimization and reduce computational cost. When the training set is small, we use an innovative data augmentation approach that can generate many novel samples of the 2D landmark points of a given set of 3D points using several camera models.
3. Compared with previous algorithms, our proposed multilayer neural network can be trained very quickly. Additionally, in testing, our algorithm runs much faster than real-time ($> 1,000$ frames/s on an i7).

2 Related Work

Directly related to our work are several approaches that reconstruct 3D shape from a single image [2], [3], [4], [5], [12]. Previous methods usually fit a shape-space model to recover the 3D configuration of the object from the 2D locations of landmarks in a single image [2], [3], [13], [14]. In [3], the authors address the human pose recovery problem by presenting an activity-independent method, given 2D locations of anatomical landmarks. Unfortunately, this method is limited to the reconstruction of human bodies. Lin *et al.* [4] derived an algorithm to recover 3D models of cars. In [13], an optimization method for location analysis is proposed to map the discriminant parts of the object into a 3D geometric representation. In [2], a shape-based model is designed, which represents a 3D shape as a linear combination of shape bases. In [14], the authors extend this previous work to deal with outliers in the 2D correspondences and rotation asynchronization. More broadly, but also related to our work, reconstructing 3D shape from a sequence of 2D images using structure from motion (SfM) has been extensively investigated [15]. In particular, the work of [6] can recover the 3D shape of an objects from 2D landmark points on a single image using SfM-trained models. 3D surface reconstruction is also related to our work, e.g., two camera models (weak-perspective and perspective) were used to reconstruct a deformable surface from a single image and a template 3D shape [16].

Except [16], the works summarized above assume a low-dimensional embedding of the underlying shape space, and represent each 3D shape as a linear combination of shape bases [17], [18]. A major limitation of linear models is that when the true (but unknown) mapping function is nonlinear, the accuracy drops significantly [19]. Methods such as [20], [21], [22], [23] fall into the category of nonlinear mapping, but like some of the aforementioned algorithms, they are limited to certain types of object categories, e.g., [3] is designed for reconstructing 3D human pose from 2D image landmarks and [4] focuses on 3D car modeling. These algorithms thus make prior assumptions or use special geometric properties of the object to constrain the solution space that do not typically generalize well to other object categories.

An additional limitation of the above cited papers is their inability to learn from either large or small datasets. Some existing algorithms require very small training sets, but are unable to yield improvements when larger datasets are available. On the other hand, some algorithms require very large training sets and are incapable of yielding reasonable results when the number of samples is small.

3 Proposed Approach

In this section, we describe how to recover the 3D shape of an object given its 2D landmarks on a single image.

3.1 Preliminaries

Let us denote the n 2D landmark points on the i^{th} image

$$\mathbf{W}_i = \begin{pmatrix} u_{i1} & u_{i2} & \cdots & u_{in} \\ v_{i1} & v_{i2} & \cdots & v_{in} \end{pmatrix} \in \mathbb{R}^{2 \times n}, \quad (1)$$

where $(u_{ij}, v_{ij})^T$ are the 2D coordinates of the landmark points in the j^{th} image. Our goal is to recover the 3D coordinates of these 2D landmark points,

$$\mathbf{S}_i = \begin{pmatrix} x_{i1} & x_{i2} & \cdots & x_{in} \\ y_{i1} & y_{i2} & \cdots & y_{in} \\ z_{i1} & z_{i2} & \cdots & z_{in} \end{pmatrix} \in \mathbb{R}^{3 \times n}, \quad (2)$$

where $(x_{ij}, y_{ij}, z_{ij})^T$ are the 3D coordinates of the landmarks in the j^{th} image.

Assuming a weak-perspective camera model, with calibrated camera matrix $\mathcal{M} = \begin{pmatrix} \lambda & 0 & 0 \\ 0 & \lambda & 0 \end{pmatrix}$, the weak-perspective projection of the object 3D landmark points is given by

$$\mathbf{W}_i = \mathcal{M} \mathbf{S}_i. \quad (3)$$

This result is of course defined up to scale, since $\mathbf{u}_i = \lambda \mathbf{x}_i$ and $\mathbf{v}_i = \lambda \mathbf{y}_i$, where $\mathbf{x}_i^T = (x_{i1}, x_{i2}, \dots, x_{in})$, $\mathbf{y}_i^T = (y_{i1}, y_{i2}, \dots, y_{in})$, $\mathbf{z}_i^T = (z_{i1}, z_{i2}, \dots, z_{in})$, $\mathbf{u}_i^T = (u_{i1}, u_{i2}, \dots, u_{in})$ and $\mathbf{v}_i^T = (v_{i1}, v_{i2}, \dots, v_{in})$. This will require that we standardize our variables when deriving our algorithm.

3.2 Deep 3D Shape Reconstruction from 2D Landmarks

3.2.1 Proposed Neural Network.—Given a training set with m 3D landmark points $\{\mathbf{S}_i\}_{i=1}^m$, we aim to learn the function $f: \mathbb{R}^{2n} \rightarrow \mathbb{R}^n$, that is,

$$\hat{\mathbf{z}}_i = f(\hat{\mathbf{x}}_i, \hat{\mathbf{y}}_i), \quad (4)$$

where $\hat{\mathbf{x}}_i$, $\hat{\mathbf{y}}_i$ and $\hat{\mathbf{z}}_i$ are obtained by standardizing $\mathbf{x}_i$, $\mathbf{y}_i$ and $\mathbf{z}_i$ to eliminate the effect of scaling and 3D translation,

$$\begin{aligned}\hat{x}_{ij} &= \frac{x_{ij} - \bar{x}_i}{(\sigma(\mathbf{x}_i) + \sigma(\mathbf{y}_i))/2}, \\ \hat{y}_{ij} &= \frac{y_{ij} - \bar{y}_i}{(\sigma(\mathbf{x}_i) + \sigma(\mathbf{y}_i))/2}, \\ \hat{z}_{ij} &= \frac{z_{ij} - \bar{z}_i}{(\sigma(\mathbf{x}_i) + \sigma(\mathbf{y}_i))/2},\end{aligned}\quad (5)$$

$\bar{x}_i$, $\bar{y}_i$ and $\bar{z}_i$ are mean values, and $\sigma(\mathbf{x}_i)$, $\sigma(\mathbf{y}_i)$ and $\sigma(\mathbf{z}_i)$ are the standard deviations.

By estimating f using a training set, we learn the geometric property of the given object. Herein, we model the function $f(\cdot)$ using a multilayer neural network.

Figure 1 depicts the overall architecture of our neural network. It contains L layers. The l^{th} layer is defined as,

$$\mathbf{a}^{(l+1)} = \tanh(\mathbf{\Omega}^{(l)} \mathbf{a}^{(l)} + \mathbf{b}^{(l)}), \quad (6)$$

where $\mathbf{a}^{(l)} \in \mathbb{R}^d$ is an input vector, $\mathbf{a}^{(l+1)} \in \mathbb{R}^r$ is the output vector, d and r specify the number of input and output nodes, respectively, and $\mathbf{\Omega} \in \mathbb{R}^{r \times d}$ and $\mathbf{b} \in \mathbb{R}^r$ are network parameters, with the former a weighting matrix and the latter a bias vector. Our neural network uses a Hyperbolic Tangent function, $\tanh(\cdot)$.

Our objective is to minimize the sum of Euclidean distances between the predicted depth location $\mathbf{a}_i^{(L)}$ and the ground-truth $\hat{\mathbf{z}}_i$ of our m 3D landmark points. Formally,

$$\min \sum_{i=1}^m \|\hat{\mathbf{z}}_i - \mathbf{a}_i^{(L)}\|_2^2, \quad (7)$$

with $\|\cdot\|$ the Euclidean metric.

We utilize the RMSProp algorithm to optimize our model parameters. In a multilayer neural network, the appropriate learning rates can vary widely during learning and between different parameters. RMSProp is a technique that updates parameters of a neural network to improve learning. It can adaptively adjust the learning rate of each parameter separately to improve convergence to a solution.

3.2.2 Testing.—When testing on the t^{th} object, we have $\mathbf{u}_t$ and $\mathbf{v}_t$ and want to estimate $\mathbf{x}_t$, $\mathbf{y}_t$ and $\mathbf{z}_t$. From Eq. (3) we have $\mathbf{u}_t = \lambda \mathbf{x}_t$ and $\mathbf{v}_t = \lambda \mathbf{y}_t$. Thus, we first standardize the data,

$$\begin{aligned}\hat{u}_{ij} &= \frac{u_{ij} - \bar{u}_t}{(\sigma(u_t) + \sigma(v_t))/2}, \\ \hat{v}_{ij} &= \frac{v_{ij} - \bar{v}_t}{(\sigma(u_t) + \sigma(v_t))/2}.\end{aligned}\quad (8)$$

This yields $\hat{x}_t = \hat{u}_t$ and $\hat{y}_t = \hat{v}_t$. This means we can directly feed $(\hat{u}_t, \hat{v}_t)$ into the trained neural network to obtain its depth $\hat{z}_t$. Then, the 3D shape of the object can be recovered as

$(\hat{u}_t^T, \hat{v}_t^T, \hat{z}_t^T)^T$, a result that is defined up to scale.

3.2.3 Missing Data—To solve the problem of missing data, we add a recurrent layer [24] on top of the multi-layer neural network to jointly estimate both the 2D coordinates of missing 2D landmarks and their depth. The complete network is trained in an end-to-end fashion, Figure 2. The module named “A” corresponds to the recurrent layer that estimates the 2D entries of the missing data, while “B” is the multi-layer neural network described before and summarized in Figure 1. The output of “A” is thus the full set of 2D landmarks and the output of “B” their corresponding depth values. The module “C” merges the outputs of “A” and “B” to generate the final output, $(\hat{u}_t^T, \hat{v}_t^T, \hat{z}_t^T)^T$, and $\mathcal{L}$ is the loss function used in this module.

In the recurrent network, we use the notation $\hat{u}_{ij}^{(s)}$ and $\hat{v}_{ij}^{(s)}$ to specify the estimated values of $\hat{u}_{ij}$ and $\hat{v}_{ij}$ at iteration s . Here, i specifies the j^{th} sample. The input to our above defined network can then be written as $\mathbf{d}_i^{(0)} = (\hat{u}_{i1}^{(0)}, \hat{v}_{i1}^{(0)}, \dots, \hat{u}_{in}^{(0)}, \hat{v}_{in}^{(0)})$, with $s = 0$ specifying the initial input. If the values of u_{ij} and v_{ij} are missing, then $\hat{u}_{ij}^{(0)}$ and $\hat{v}_{ij}^{(0)}$ are set to zero. Otherwise the values of u_{ij} and v_{ij} are standardized using Eq. (8) to obtain $\hat{u}_{ij}^{(0)}$ and $\hat{v}_{ij}^{(0)}$.

In subsequent iterations, from $s-1$ to s , if the j^{th} landmark is not missing, $\hat{u}_{ij}^{(s)} = \hat{u}_{ij}^{(s-1)}$ and $\hat{v}_{ij}^{(s)} = \hat{v}_{ij}^{(s-1)}$. If the j^{th} landmark is missing, then $\hat{u}_{ij}^{(s)} = g(\sum_{k=1}^{2n} w_{k(2j-1)} d_{ik}^{(s-1)})$, $\hat{v}_{ij}^{(s)} = g(\sum_{k=1}^{2n} w_{k(2j)} d_{ik}^{(s-1)})$, where $g(\cdot)$ can be the identity function or a nonlinear function (e.g. $\tanh(\cdot)$) and $w_{k(2j-1)}$, $w_{k(2j)}$, $k = 1, \dots, 2n$, $j = 1, \dots, n$ are the parameters of the recurrent layer. In our experiments we find that $g(\cdot)$ being the identity function works best.

We set the number of iterations to τ , which yields $\mathbf{d}_i = \sum_{s=1}^{\tau} \lambda_s \mathbf{d}_i^{(s)}$ as the final output of the recurrent layer, where λ_s are learned weights. We initialize λ_s as $0 < \lambda_1 < \dots < \lambda_\tau$ and $\sum_{s=1}^{\tau} \lambda_s = 1$. The vector $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_\tau)^T$ is then learned using backpropagation. By using the weighted sum of the output at each step rather than the output at the last step as final output of the recurrent layer, we can enforce intermediate supervision to make the recurrent layer gradually converge to the desirable output.

3.2.4 Data augmentation approach.—In many applications the number of training samples (i.e., 2D and corresponding 3D landmark points) is small. However, any regressor designed to learn the mapping function $f(\cdot)$ requires of a large number of training samples with the 2D landmarks as seen from as many cameras and views (i.e., translation, rotation) as possible. We resolve this with a simple, yet effective data augmentation approach.

The key to our approach is to note that, for a given object, its 3D structure does not change. What changes are the 2D coordinates of the landmark points in the image. For example, scaling or rotating an object in 3D yields different 2D coordinates of the same object landmarks. Thus, our task is to generate as many of these possible sample views as possible. We do this with a camera model.

Herein, we use an affine camera model to generate a very large number of images of the known 3D sample objects. We do this by varying the extrinsic parameters of the camera model (e.g., rotation). Specifically, we use the weak-perspective camera model defined above.

We also use this data augmentation step to model imprecisely localized 2D landmark points. All detection algorithms yield imprecise detections (even when fiducial detections are done by humans) [25], [26]. We address this problem by modeling the detection error as Gaussian noise, with zero mean and variance σ . Specifically, we use a small variance equivalent to about 3% the size of the object. This means that, in addition to the 2D landmark points given by the camera models used above, we will incorporate 2D landmark points that have been altered by adding this random Gaussian noise. This allows our neural network to learn to accurately recover the 3D shape of an object from imprecisely localized 2D landmark points.

It is important to note that, when the original training set is small, we can still train our neural network effectively using this trick. In fact, we have found experimentally that we do not need a large number of training samples to obtain extremely low reconstruction errors using our derived approach and this data augmentation trick. *This is significant because deep neural nets usually require very large training sets to learn to detect and recognize objects in images.*

Of course, even when the number of samples is large, our approach helps reduce the 3D reconstruction error by incorporating intrinsic and extrinsic camera parameters and detection errors not well represented in the samples.

3.2.5 Implementation Details.—Our feed-forward neural network contains six layers. The number of nodes in each layer is $[2n, 2n, 2n, 2n, 2n, n]$. We divide our training data into a training and a validation set. In each of these two sets, we perform data augmentation.

We use Keras library [27] on top of Theano [28] to implement our proposed multilayer neural network. Early stopping is enabled to prevent overfitting. We stop the training process if the validation error does not decrease after 10 iterations. We set the initial learning rate at .01.

4 Experimental Results

We provide experiments on a variety of databases to test the effectiveness of our algorithm. We used the following datasets: the Human 3.6M database [29], the CMU Motion Capture database [30], the fine-grained 3D car (FG3DCar) database [4], the Binghamton-Pittsburgh 4D Spontaneous Expression (BU-3DFE) database [31], [32] and the flag flapping in the wind database [33]. We also report our results on the sequestered dataset of the 3D Face Alignment in the Wild Challenge (3DFAW), done in conjunction with the 2016 European Conference on Computer Vision (ECCV), where the herein derived algorithm was a top performer.

Comparisons with the state-of-the-art demonstrates that our algorithm is significantly more accurate and efficient in recovering 3D shape from a single 2D image than previously defined methods. The 3D reconstruction error is evaluated using the Procrustes distance between the reconstructed shape and the ground-truth. The ground-truth is given in millimeters (mm) in the original dataset. But, the ground-truth is normalized to a standard scale (given by Eq. (5)) to make the error measure invariant to scale.

Also, to demonstrate the robustness of our method, we perform sensitivity analysis on these databases. Results show that our method is tolerant to moderate Gaussian noise. We also show results of how our method handles missing data.

4.1 Human 3.6M Database

The Human 3.6M database [29] contains 3.6 million 3D human poses and corresponding images from 17 scenarios such as discussion and smoking. The motions in the database were performed by 11 professional actors (5 female and 6 male).

As suggested by [35], the data from five subjects (S1, S5, S6, S7, S8) are used for training, and the data from two subjects (S9, S11) are used for testing. During testing, we also truncate the frames beyond 30 seconds for each sequence. The frame rate is downsampled from original 50 fps to 10 fps. Each body shape contains 32 joints. Unlike the paper [35], which used video sequence to do the training and testing, we only use image information and treat each frame independently. The comparison results are shown in Table 1. Our result is even better than the method which requires a sequence to train and test.

4.2 CMU Motion Capture Database

The CMU motion capture database contains 3D human body joints locations of subjects performing various physical activities. It contains 144 subjects and the following categories: human interaction, interaction with environment (playground, uneven terrain and path with obstacles), locomotion (running, walking, jumping and varied), physical activities & sports (basketball, dance, gymnastics, acrobatics, martial arts, racquet/paddle sports, soccer, boxing, general exercise and stretching, golf and frisbee), situations & scenarios (common behaviors and expressions, pantomime, communication gestures and signals and cross-category variety) and test motions.

Each body shape is defined by 15 3D landmark points. To provide a fair comparison, we follow the experimental setting of [2], where sequences of subject 86 are used for training and sequences of subjects 13, 14 and 15 are used for testing. We train the neural network in 10,000 epochs. We train one batch for at most 300 iterations within one epoch. The batch is either the original training data (the first epoch) or a random rotation of the original training data in 3D space. To represent real human shape in a 2D image, we set the rotation angle to be uniformly distributed within the range of $[-20^\circ, 20^\circ]$ about the x axis, $[-20^\circ, 20^\circ]$ about the y axis and $[-180^\circ, 180^\circ]$ about the z axis.

Comparative results are in Table 2. As seen in this table, our results are significantly more accurate than those given by previous algorithms. To further demonstrate the effectiveness and generality of our method, we randomly selected several 2D images from the human-in-the-wild dataset [34] and used the herein derived algorithm to recover the 3D shape of the human bodies in these images. The results are in Figure 3.

To further verify that our proposed approach is better than state-of-the-art algorithms, we compare our method with the one proposed in [36]. We directly test the public available trained model [36] on the Joint-limit dataset proposed in this paper. The testing error is 0.039. We train and test our proposed model on the Joint-limit dataset which yields a testing error of 0.035. Our method achieves better performance compared with this pose-specific method.

4.3 3D Face Reconstruction

The BU-3DFE database has images of 100 subjects performing six facial expressions in front of a 3D scanner. Each subject performed each expression 25 times, for a total of 2,500 3D sample images. Every sample has 83 annotated 3D landmarks. We randomly select 60 subjects for training, 10 for validation and the remaining 30 for testing.

We trained our neural network in 50,000 epochs and used the same training strategy as described in the preceding section, with the difference that we restricted rotations about the z axis in the range of $[-60^\circ, 60^\circ]$ to more accurately reflect real head movement.

The mean testing error of our algorithm is **0.0032**. To provide comparative results against the state-of-the-art, we train and test the algorithm of [2] and use the same experimental setting described above. Comparative results are in Table 3.

We also validate the cross database generality of our trained model on the Helen database of [37]. Here, we randomly select some face images and manually label the face landmarks. The reconstructed 3D shapes are shown in Figure 4.

To provide an additional unbiased result, we participated in the 3DFAW competition¹, which was conducted in conjunction with ECCV. Three of the four datasets in the challenge are subsets of MultiPIE [38], BU-4DFE [39] and BP4D-Spontaneous [40] databases respectively. Another dataset TimeSlice3D contains annotated 2D images that are extracted

¹<https://competitions.codalab.org/competitions/10261>

from online videos. In total, there are 18,694 images. Each image has 66 labeled 3D fiducial points and a face bounding box centered around the mean 2D projection of the landmarks. The 2D to 3D correspondence assumes a weak-perspective projection. The depth values have been normalized to have zero mean. Because this competition required that we also estimate the 2D landmark points in the image, we incorporated an additional deep network to our architecture to the model summarized in Figure 1 and trained it to detect 2D face landmarks [41]. Details of this deep net are in the Supplementary Documentation.

Reconstruction error was reported by the organizers of the competition on a sequestered dataset to which we did not have prior access. Our algorithm was a top performer, with a significant margin over other methods. Herein, we also provide comparative results against the algorithm of [2] in Table 3.

4.4 FG3DCar Database

This is a fine-grained 3D car database. It consists of 300 images of 30 different car models as seen from multiple views. Each car model has 10 images. Each image has 64 annotated 2D landmarks and a 3D shape reconstructed CAD (Computer-Aided Design) model. We adopt the default setting for training and testing, i.e., half of the 3D shapes of each car model is used for training, the other half is used for testing. In order to train our neural network, we further split the 150 3D training sample shapes into training (120 shapes) and validation (30 shapes) sets. Testing is conducted on the remaining 150 images.

The neural network is trained for 100,000 epochs. We follow the same procedure used in the preceding two sections and train our neural network on one batch for at most 300 iterations in each epoch. The batch is either composed of the 120 training samples (the first epoch) or a random 3D rotation of the 120 training samples. We augment the validation set 2,000 times using the data augmentation approach described above, resulting in a total of 60,000 validation sample images.

The mean 3D reconstruction error on the testing set is **0.0020**. Comparative results with the method of [2] are in Table 3. Qualitative examples of our results are shown in Figure 5.

4.5 Flag Flapping in the Wind Database

The flag flapping in the wind database [33] has 450 frames, each of which has 540 vertices. We use the first 300 frames for training and the rest for testing. The network is trained in 30,000 epochs and we use the same procedure described in the preceding sections.

The mean testing error is **0.0004**. Figure 6 shows some of our reconstructed results compared with the ground-truth. The reconstructed 3D shape is shown using filled red circles and the 3D ground-truth with open green circles. As we can see in these results, the reconstructed and true shape are almost identical.

4.6 Discussion on training strategy

As we saw in the results above, the proposed algorithm yields significantly superior results to those given by state-of-the-art algorithms. In this section we evaluate the contribution of the proposed data augmentation approach. To evaluate this contribution, we show

reconstruction errors obtained with and without our data augmentation approach in Table 4. As can be seen in the table, the proposed data augmentation approach has a significant contribution on the final results.

In this table, we also show the results we obtained when using the whole training dataset or 3/4, 1/2 and 1/4 of the dataset. As we can see in the table, our data augmentation approach allows us to train our algorithm with as little as 1/4th of the dataset and still obtained comparable results to those attained when we utilize all the training set.

4.7 Noise and missing data

To determine how sensitive the proposed neural network is to inaccurate 2D landmark detections, we add independent random Gaussian noise with variance σ to the databases used in the preceding sections. That is, we add noise to the training and testing samples. Specifically, we apply Gaussian noise to the 2D landmarks.

Figure 7(a) shows how little the performance degrades as σ increases when noise is added to the CMU Motion Capture database. The average height of subjects in this dataset is 1,500 mm, meaning the variance of the noise σ is about 3%. In Figure 7(a), we can see the robustness of the proposed algorithm to inaccurate 2D landmark positions. Figure 7(e) shows the relative reconstruction error averaged across the testing subjects for each landmark with and without noise.

The results on the BU-3DFE Face Database, FG3DCar Database and Flag Flapping in the Wind sequence are shown in Figures 7(b)-7(d) and 7(f)-7(h). The average width of the faces in BU-3DFE is 140 mm, hence, the variance is 5%. The mean width of the car models in FG3DCar is 569 pixel, hence, the variance is 2%. The mean width of the flags is 386 mm, hence, the variance is 3%.

Finally, we tested the ability of the trained system to deal with missing data. Here, each training and validation sample had either 2.5% or 5% randomly selected landmark points missing; during training *and* testing. Comparative results with different proportions of missing landmark points are in Table 5. As can be seen in the table, our paper can achieve smaller reconstruction errors compared with [2], [35], even when ours has missing data and that of [2], [35] do not. We also compare our algorithm with a simple nearest neighbor approach. In this approach, for each testing sample, its 3D reconstruction is the 3D shape in the training set whose 2D projection has the smallest Euclidean distance with that of the test image. Figure 8 shows qualitative results on a few images of humans in the wild and face in the wild databases, respectively.

5 Conclusions

We have presented a simple algorithm for the reconstruction of 3D shapes from 2D landmark points that yields extremely low reconstruction errors. Specifically, we proposed to use a feed-forward neural network to learn the mapping function between a set of 2D landmark points and an object's 3D shape. The exact same neural network is used to learn the mappings of rigid (e.g., cars), articulated (e.g., human bodies), nonrigid (e.g., faces), and

highly-deformable objects (e.g., flags). The system performs extremely well in all cases and yields results as much as two-fold better than previous state-of-the-art algorithms. This neural network runs much faster than real-time, > 1000 frames/s, and can be trained with small sample sets.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgment

This research was supported in part by the National Institutes of Health, grant R01-DC-014498, a Human Frontier Science Program, grant RGP0036, and by a Google Faculty Research Award to AMM. We thank the reviewers for their constructive feedback.

References

- [1]. Martinez AM, "Computational models of face perception," Current Directions in Psychological Science, 2017.
- [2]. Zhou X, Leonardos S, Hu X, and Daniilidis K, "3d shape estimation from 2d landmarks: A convex relaxation approach," in The IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2015.
- [3]. Ramakrishna V, Kanade T, and Sheikh Y, "Reconstructing 3d human pose from 2d image landmarks," in ECCV, 2012, pp. 573–586.
- [4]. Lin Y-L, Morariu VI, Hsu W, and Davis LS, "Jointly optimizing 3d model fitting and fine-grained classification," in ECCV. Springer, 2014, pp. 466–480.
- [5]. Kar A, Tulsiani S, Carreira J, and Malik J, "Category-specific object reconstruction from a single image," in The IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2015.
- [6]. Hamsici OC, Gotardo PFU, and Martinez AM, "Learning spatially-smooth mappings in non-rigid structure from motion," in ECCV, 2012, pp. 260–273.
- [7]. Fayad J, Russell C, and Agapito L, "Automated articulated structure and 3d shape recovery from point correspondences," in Computer Vision (ICCV), 2011 IEEE International Conference on. IEEE, 2011, pp. 431–438.
- [8]. Gotardo PFU and Martinez AM, "Kernel non-rigid structure from motion," in Computer Vision (ICCV), 2011 IEEE International Conference on, 2011, pp. 802–809.
- [9]. Jeni LA, Cohn JF, and Kanade T, "Dense 3d face alignment from 2d videos in real-time," in Automatic Face and Gesture Recognition (FG), 2015 11th IEEE International Conference and Workshops on, vol. 1 IEEE, 2015, pp. 1–8.
- [10]. Russell C, Fayad J, and Agapito L, "Energy based multiple model fitting for non-rigid structure from motion," in Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on IEEE, 2011, pp. 3009–3016.
- [11]. Russell C, Yu R, and Agapito L, "Video pop-up: Monocular 3d reconstruction of dynamic scenes," in European Conference on Computer Vision Springer, 2014, pp. 583–598.
- [12]. Vicente S, Carreira J, Agapito L, and Batista J, "Reconstructing pascal voc," in Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on IEEE, 2014, pp. 41–48.
- [13]. Zhu M, Zhou X, and Daniilidis K, "Single image pop-up from discriminatively learned parts," in The IEEE International Conference on Computer Vision (ICCV), 2015.
- [14]. Zhou X, Zhu M, Leonardos S, and Daniilidis K, "Sparse representation for 3d shape estimation: A convex relaxation approach," arXiv preprint arXiv: 1509.04309, 2015.
- [15]. Hartley R and Zisserman A, Multiple view geometry in computer vision. Cambridge university press, 2003.

- [16]. Bartoli A and Collins T, "Template-based isometric deformable 3d reconstruction with sampling-based focal length self-calibration," in 2013 IEEE Conference on Computer Vision and Pattern Recognition, 2013, pp. 1514–1521.
- [17]. Gotardo PF and Martinez AM, "Computing smooth time trajectories for camera and deformable shape in structure from motion with occlusion," IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 33, no. 10, pp. 2051–2065, 2011. [PubMed: 21383398]
- [18]. Akhter I, Sheikh Y, Khan S, and Kanade T, "Trajectory space: A dual representation for nonrigid structure from motion," Pattern Analysis and Machine Intelligence, IEEE Transactions on, vol. 33, no. 7, pp. 1442–1456, 2011.
- [19]. Gotardo PF and Martinez AM, "Kernel non-rigid structure from motion," in 2011 International Conference on Computer Vision IEEE, 2011, pp. 802–809.
- [20]. Urtasun R, Fleet DJ, Hertzmann A, and Fua P, "Priors for people tracking from small training sets," in Tenth IEEE International Conference on Computer Vision (ICCV'05) Volume 1, vol. 1, 2005, pp. 403–410 Vol. 1.
- [21]. Salzmann M, Urtasun R, and Fua P, "Local deformation models for monocular 3d shape recovery," in 2008 IEEE Conference on Computer Vision and Pattern Recognition, 2008, pp. 1–8.
- [22]. Dame A, Prisacariu VA, Ren CY, and Reid I, "Dense reconstruction using 3d object shape priors," in 2013 IEEE Conference on Computer Vision and Pattern Recognition, 2013, pp. 1288–1295.
- [23]. Chen Y, Kim T-K, and Cipolla R, Inferring 3D Shapes and Deformations from Single Views, 2010, pp. 300–313.
- [24]. Bengio Y and Gingras F, "Recurrent neural networks for missing or asynchronous data," Advances in neural information processing systems, 1996.
- [25]. Ding L and Martinez AM, "Features versus context: An approach for precise and detailed detection and delineation of faces and facial features," IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 32, no. 11, pp. 2022–2038, 2010. [PubMed: 20847391]
- [26]. Martinez AM, "Recognizing imprecisely localized, partially occluded, and expression variant faces from a single sample per class," IEEE Transactions on Pattern analysis and machine intelligence, vol. 24, no. 6, pp. 748–763, 2002.
- [27]. Chollet F, "keras," <https://github.com/fchollet/keras>, 2015.
- [28]. Bastien F, Lamblin P, Pascanu R, Bergstra J, Goodfellow I, Bergeron A, Bouchard N, Warde-Farley D, and Bengio Y, "Theano: new features and speed improvements," arXiv preprint arXiv: 1211.5590, 2012.
- [29]. Ionescu C, Papava D, Olaru V, and Sminchisescu C, "Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments," IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 36, no. 7, pp. 1325–1339, 7 2014. [PubMed: 26353306]
- [30]. MoCap. Carnegie mellon university graphics lab motion capture database. <http://mocap.cs.cmu.edu/>.
- [31]. Zhang X, Yin L, Cohn JF, Canavan S, Reale M, Horowitz A, Liu P, and Girard JM, "Bp4d-spontaneous: a high-resolution spontaneous 3d dynamic facial expression database," Image and Vision Computing, vol. 32, no. 10, pp. 692–706, 2014.
- [32]. Zhang X, Yin L, Cohn JF, Canavan S, Reale M, Horowitz A, and Liu P, "A high-resolution spontaneous 3d dynamic facial expression database," in Automatic Face and Gesture Recognition (FG), 2013 10th IEEE International Conference and Workshops on, 2013, pp. 1–6.
- [33]. White R, Crane K, and Forsyth DA, "Capturing and animating occluded cloth," in *ACM Transactions on Graphics (TOG)*, vol. 26, no. 3, 2007, p. 34.
- [34]. Ramanan D, "Learning to parse images of articulated bodies," in Advances in Neural Information Processing Systems 19, Schölkopf B, Platt JC, and Hoffman T, Eds., 2007, pp. 1129–1136.
- [35]. Zhou X, Zhu M, Leonardos S, Derpanis KG, and Daniilidis K, "Sparseness meets deepness: 3d human pose estimation from monocular video," in The IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 6 2016.

- [36]. Akhter I and Black MJ, "Pose-conditioned joint angle limits for 3D human pose reconstruction," in IEEE Conf. on Computer Vision and Pattern Recognition (CVPR 2015), Jun. 2015, pp. 1446–1455.
- [37]. Le V, Brandt J, Lin Z, Bourdev L, and Huang TS, "Interactive facial feature localization," in Proceedings of the 12th European Conference on Computer Vision - Volume Part III, ser. ECCV'12, 2012, pp. 679–692.
- [38]. Gross R, Matthews I, Cohn J, Kanade T, and Baker S, "Multi-pie," Image and Vision Computing, vol. 28, no. 5, pp. 807–813, 2010.
- [39]. Yin L, Chen X, Sun Y, Worm T, and Reale M, "A highresolution 3d dynamic facial expression database," in Automatic Face & Gesture Recognition, 2008. FG'08. 8th IEEE International Conference On IEEE, 2008, pp. 1–6.
- [40]. Zhang X, Yin L, Cohn JF, Canavan S, Reale M, Horowitz A, Liu P, and Girard JM, "Bp4d-spontaneous: a high-resolution spontaneous 3d dynamic facial expression database," Image and Vision Computing, vol. 32, no. 10, pp. 692–706, 2014.
- [41]. Zhao R, Benitez-Quiroz F, Wang Y, Liu Y, and Martinez A, "Fast and precise face alignment and 3d shape reconstruction from a single 2d image," in Proceedings of the 14th European Conference on Computer Vision Workshop, ser. ECCVW'16, 2016.

Fig. 1:

Framework of the proposed algorithm to estimate the 3D shape of a set of 2D landmark points on a single image. The resulting 3D shape is estimated in real-time ($> 1,000$ frames/s on an i7 desktop).

Fig. 2:

Deep network for dealing with missing data. $d^{(0)}$ is the input to the network, “A” is the recurrent layer with τ steps for estimating missing inputs. “B” is the multi-layer neural network defined in Section 3.2.1 and summarized in Figure 1. “C” combines the results of “A” and “B” to yield the final output of the network.

Fig. 3:

Results of our algorithm on images of the humans in-the-wild dataset [34]. The first and fifth columns are the single 2D images and 2D landmarks (in yellow) used by our algorithm. The reconstructed 3D shapes are shown on the second, third, fourth, sixth, seventh and eighth columns; these show the 3D reconstruction from multiple views to more clearly demonstrate the quality of the recovered 3D shape.

Fig. 4:

Illustration of our 3D estimation results on the Helen database [37]. The model is trained on the BU-3DFE face database and tested on Helen. The first and fifth columns show the input 2D images with landmarks (in yellow); the other columns show the estimated 3D shapes.

Fig. 5: Illustration of our 3D estimation results on the FG3DCar database. The first and fifth columns corresponds to the input 2D images and landmark points (in yellow), with the other columns showing the recovered 3D shapes.

Fig. 6:
3D shape reconstruction results on the flag flapping in the wind sequence. Comparisons of the reconstructed 3D shape (in red) with its 3D ground-truth (in green).

Fig. 7:

Additive random Gaussian noise is used to test the robustness of the proposed algorithm to 2D inaccurate detections on CMU Motion Capture Database (a), BU-3DFE Database (b), FG3DCar Database (c) and Flag Flapping in the Wind Database (d). Reconstruction error is shown on the y axis and σ on the x axis. Note the tiny reconstruction errors even for large values of σ . Sensitivity of reconstruction to each landmark when σ increases on CMU Motion Capture Database (e), BU-3DFE Database (f), FG3DCar Database (g) and Flag Flapping in the Wind Database (h). The radius of the circle indicates the relative reconstruction error for each landmarks.

Fig. 8:

Qualitative results. Top row: Missing landmarks are identified with a dotted line in the images. Bottom row: Missing landmarks are marked in red in the images.

TABLE 1:

Comparative 3D mean reconstruction errors of our method and Zhou *et al.* [35] on the Human 3.6M database.

This paper	Zhou <i>et al.</i> [35]
0.0120	0.0359

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

TABLE 2:

Comparisons of the reconstruction error on CMU Motion Capture Database.

Methods	Subject 13	Subject 14	Subject 15
This paper	.0231	.0200	.0095
Zhou <i>et al.</i> [2]	.0653	.0643	.0405
Ramakrishna <i>et al.</i> [3]	.0983	.0979	.0675

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

TABLE 3:

Comparative results of our approach and Zhou *et al.* [2] on BU-3DFE, 3DFAW and FG3DCar.

Database	This paper	Zhou <i>et al.</i> [2]
BU-3DFE Face	0.0032	0.0053
3DFAW	0.0042	0.0055
FG3DCar	0.0020	0.0042

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

TABLE 4:

Results with different number of training samples, w/ and w/o data augmentation on BU-3DFE.

Training Strategy	Error
whole training samples, w/o data augmentation	0.0169
1/4 training samples, w/ data augmentation	0.0036
1/2 training samples, w/ data augmentation	0.0034
3/4 training samples, w/ data augmentation	0.0033
whole training samples, w/ data augmentation	0.0032

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

TABLE 5:

Comparisons of the reconstruction error on five databases with different proportions of missing data entries.

Methods	Human 3.6M	CMU MoCap			BU-3DFE Face	FG3DCar	Flag Flapping
		Subject I3	Subject I4	Subject I5			
This paper	0.0120	0.0231	0.0200	0.0095	0.0032	0.0020	0.0004
This paper (w/ 2.5% missing)	0.0314	0.0413	0.0396	0.0307	0.0037	0.0086	0.0065
This paper (w/ 5% missing)	0.0383	0.0728	0.0694	0.0693	0.0053	0.0086	0.0076
Nearest neighbor (w/ 2.5% missing)	0.0426	0.0983	0.0844	0.0497	0.0106	0.0123	0.0101
Nearest neighbor (w/ 5% missing)	0.0428	0.0992	0.0859	0.0509	0.0107	0.0124	0.0102
Zhou <i>et al.</i> [2]	–	0.0653	0.0643	0.0405	0.0053	0.0042	–
Zhou <i>et al.</i> [35]	0.0359	–	–	–	–	–	–
Ramakrishna <i>et al.</i> [3]	–	0.0983	0.0979	0.0675	–	–	–