

# Multi-view 3D Models from Single Images with a Convolutional Network

Maxim Tatarchenko, Alexey Dosovitskiy, Thomas Brox

Department of Computer Science  
University of Freiburg  
{tatarchm, dosovits, brox}@cs.uni-freiburg.de

**Abstract.** We present a convolutional network capable of inferring a 3D representation of a previously unseen object given a single image of this object. Concretely, the network can predict an RGB image and a depth map of the object as seen from an arbitrary view. Several of these depth maps fused together give a full point cloud of the object. The point cloud can in turn be transformed into a surface mesh. The network is trained on renderings of synthetic 3D models of cars and chairs. It successfully deals with objects on cluttered background and generates reasonable predictions for real images of cars.


**Keywords:** 3D from single image, deep learning, convolutional networks

## 1 Introduction

The ability to infer a 3D model of an object from a single image is necessary for human-level scene understanding. Despite the large success of deep learning in computer vision and the diversity of tasks being approached, 3D representations are not yet in the focus of deep networks. Can we make deep networks learn such 3D representations?

In this paper, we present a simple and elegant encoder-decoder network that infers a 3D model of an object from a single image of this object, see Figure 1. We represent the object by what we call "multi-view 3D model" – the set of all its views and corresponding depth maps. Given an arbitrary viewpoint, the network we propose generates an RGB image of the object and the depth map. This representation contains rich information about the 3D geometry of the object, but allows for more efficient implementation than voxel-based 3D models. By fusing several views from our multi-view representation we get a full 3D point cloud of the object, including parts invisible in the original input image.

While technically the task comes with many ambiguities, humans are known to be good in using their prior knowledge about similar objects to guess the missing information. The same is achieved by the proposed network: when the input image does not allow the network to infer the parts of an object – for example, because the input only shows the front view of a car and there is no



**Fig. 1.** Our network infers an object’s 3D representation from a single input image. It then predicts unseen views of this object and their depth maps. Multiple such views are fused into a full 3D point cloud, which is further optimized to obtain a mesh.

information about its back – it fantasizes the most probable shape consistent with the presented data (for example, a standard sedan car).

The network is trained end-to-end on renderings of 3D models from the ShapeNet dataset [1]. We render images on the fly during network training, with random viewpoints and lighting. This makes the training set very diverse, thanks to the size of ShapeNet, and effectively infinite. We make the task more challenging and realistic by pasting the object renderings on top of random background images. In this setup, the network learns to automatically segment out the object. Moreover, we show that networks trained on synthetic images of this kind yield reasonable predictions for real-world images without any additional adaptation.

**Contributions.** First, we largely improve on the visual quality of the generated images compared to previous work. Second, we achieve this with a simpler and thus more elegant architecture. Finally, we are the first who can apply the network to images with non-homogeneous background and natural images.

## 2 Related work

*Unseen view prediction* Our work is related to research on modeling image transformations with neural-network-based approaches. These often involve multiplicative interactions, for example gated RBMs [2], gated autoencoder [3] or Disentangling Boltzmann Machines [4]. These approaches typically do not scale to large images, although they potentially could by making use of architectures similar to convolutional DBNs [5]. They are also typically only applicable to small transformations.

Transforming autoencoders [6] are trained to generate a transformed version of an input image given the desired transformation. When applied to the NORB dataset of  $96 \times 96$  pixel stereo image pairs of objects, this approach can apply small rotations to the input image.

The multi-view perceptron [7] is a network that takes a face image and a random vector as input and generates a random view of this face together with the corresponding viewpoint. In contrast, our model can generate directly the desired view without the need for random sampling. Kulkarni et al. [8] trained a variant of a variational autoencoder with factored hidden representations, where

certain dimensions are constrained to correspond to specific factors of variations in the input data, such as viewpoint and lighting. This method is conceptually interesting and it allows to generate previously unseen views of objects, but the quality of predictions made by our network is significantly better, as we show in the experimental section.

A simplified version of unseen view prediction is predicting HOG descriptors [9] instead of images. Chen et al. [10] pose the problem as tensor completion. Su et al. [11] find object parts similar to those of a given object in a large dataset of 3D models and interpolate between the desired views of these. These methods do not learn a 3D representation of the object class but approximate unseen views by linear combinations of models from a fixed dataset.

Dosovitskiy et al. [12] trained an 'up-convolutional' network to generate an image of a chair given the chair type and a viewpoint. This method is restricted to generating images of objects from the training set or interpolating between them. Applying the method to a new test image requires re-training the network, which takes several days. While the decoder part of our network is similar to the architecture of Dosovitskiy et al., our network also includes an encoder part which infers the high-level representation from a given input image. Hence, at test time we can generate unseen views and depth maps of new objects by simply forward propagating an image of the object through the network. Our approach also yields more accurate predictions.

Most closely related is the concurrent work by Yang et al. [13, 14]. They train a recurrent network that can rotate the object in the input image: given an image, it generates a view from a viewpoint differing by a fixed increment. This makes the approach restricted to generating a discrete set of views, while we are able to vary the angle continuously. In the approach of Yang et al., one might train the network with a small angle increment and predict views at finer quantization levels than the 15 degrees used by the authors. However, this would require more recurrent iterations for performing large rotations. It would be slow and probably would lead to error accumulation. Our network does not have such restrictions and produces an arbitrary output view in a single forward pass. Moreover, it can generate a full 3D point cloud, can deal with non-homogeneous background, and the generated images are of much better quality.

*3D from single image* Inferring a 3D model of an object from a single image is a long-standing, very difficult task in computer vision. A general approach is to use certain models of lighting, reflectance and object properties to disentangle these factors given a 2D input image [15]. When reconstructing a specific object class, prior knowledge can be exploited. For example, morphable 3D models [16], [17] are commonly used for faces. Kar et al. [18] extended this concept to object categories with more variation, such as cars and chairs, and combined it with shape-from-shading to retrieve also the high frequency components of the shape. For building their morphable 3D model they rely on ideas from Vicente et al. [19], who showed that the coarse 3D structure can be reconstructed from multiple images of the same object class (but different object instances) and some keypoint annotation. In contrast to Kar et al. [18], our approach does not use

an explicit 3D model. A 3D model representation for the object class is rather implicit in the weights of the convolutional network.

Aubry et al. [20] proposed an approach for aligning 3D models of objects with images of these objects. The method makes use of discriminative part detectors and works on complicated real scenes. On the downside, this is a nearest-neighbor kind of method: it selects the best fitting 3D models from a fixed set of models. This limits the generalization capability of the method and makes it proportionally slower if the model collection grows in size.

Huang et al. [21] reconstruct 3D models from single images of objects by jointly analyzing large collections of images and 3D models of objects of the same kind. The method yields impressive results. However, it jointly processes large collections of images and models with a nearest neighbor approach and hence cannot be applied to a new image at test time that is different from all models in the dataset.

Eigen et al. [22] trained convolutional networks to predict depth from single images of indoor scenes. This is very different from our work in that we predict depth maps not only for the current viewpoint, but also for all other viewpoints. Wu et al. [23] trained 3D Convolutional Deep Belief Networks capable of generating a volumetric representation of an object from a single depth map. This method requires a depth map as input, while our networks only take a single RGB image.

### 3 Model description

We train a network that receives an input pair  $(x_i, \theta_i)$ , where  $x_i$  is the input image and  $\theta_i$  the desired viewpoint, and aims to estimate a pair  $(y_i, d_i)$ , where  $y_i$  is the 2D projection of the same object from the requested viewpoint and  $d_i$  is the depth map of this projection. While the input images  $x_i$  may have complicated background, the targets  $y_i$  always have monotonous background.  $\theta_i$  is a vector defining the viewpoint; it consists of two angles – azimuth  $\theta_i^{az}$  and elevation  $\theta_i^{el}$  – and the distance  $r$  from the object center. Angles are given by their sine and cosine to deal with periodicity. The viewpoint of the input image is *not* given to the network. This makes the task more difficult since the network must implicitly infer the viewpoint from the input image.

The network is trained by minimizing the loss function  $\mathcal{L}$  which is a weighted sum of two terms: squared Euclidean loss for the RGB image and  $L_1$  loss for the depth image:


$$\mathcal{L} = \sum_i \|y_i - \hat{y}_i\|_2^2 + \lambda \|d_i - \hat{d}_i\|_1, \quad (1)$$

where  $\hat{y}_i$  and  $\hat{d}_i$  are the outputs of the network and  $\lambda$  is the weighting coefficient. We used  $\lambda = 0.1$  in our experiments.

#### 3.1 Architecture

The architecture of our encoder-decoder network is shown in Figure 2. It is simple and elegant. The encoder part (blue in the figure) processes the input image to





**Fig. 2.** The architecture of our network. The encoder (blue) turns an input image into an abstract 3D representation. The decoder (green) processes the angle, modifies the encoded hidden representation accordingly, and renders the final image together with the depth map.

obtain a hidden 3D representation  $z_{obj}$  of an object shown in the image. The decoder part (green in the figure) then takes  $z_{obj}$  and the desired viewpoint as inputs and renders the final output image.

During training, the network is always presented with pairs of images showing two views of the same object together with the viewpoint of the output view. Objects are randomly sampled from a large database of 3D models, and pairs of views are randomly selected.

Technically, the encoder part propagates an input image through a standard ConvNet architecture, which consists of 5 convolutional layers with stride  $s = 2$  in each layer and one fully connected layer in the end. The decoder part independently processes the angle in 3 fully connected (FC) layers, then merges the resulting code with the output of the encoder and performs joint processing in 3 more FC layers. Finally, it renders the desired picture using 5 up-convolutional layers (also known as "deconvolutional"). We experimented with deeper and wider networks, but did not observe a significant difference in performance.

The up-convolutional layers perform upsampling+convolution, opposite to the standard convolution+pooling. During upsampling, each pixel is replaced with a  $2 \times 2$  block containing the original pixel value in the top left corner and zeros everywhere else. For both convolutional and up-convolutional layers of the network we use  $5 \times 5$  filters for outer layers and  $3 \times 3$  filters for deeper layers.

The Leaky ReLU nonlinearity with the negative slope 0.2 is used after all layers, except for the last one, which is followed by the *tanh*.

### 3.2 Multi-view 3D to point cloud and mesh

The multi-view 3D model provided by the network allows us to generate a point cloud representing the object, which in turn can be transformed into a mesh. To achieve this, for a single input we generate multiple output images from different viewpoints together with their corresponding depth maps. The camera parameters are known: both internal (focal length, camera model) and external



**Fig. 3.** Train-test split of cars. Sample renderings and their nearest neighbors are shown. Each row shows on the left a rendering of a query model from the test set together with several HOG space nearest neighbors from the training set. The two query models on the right are ‘difficult’ ones.

(camera pose). This allows us to reproject each depth map to a common 3D space and obtain a single point cloud.

As a post-processing step we can turn the point cloud into a dense surface model with the method of Pock et al. [24]. This method uses depth information together with the point normals to compute the final mesh. As the normal information is missing in our case (although it potentially could also be estimated by the network), we approximate it by providing the direction to the camera for each point. Since the normals are optimized anyway by the fusion method, this approximation yields good results in practice.

### 3.3 Dataset

We used synthetic data from the ShapeNet dataset [1] for training the networks. The dataset contains a large number of 3D models of objects belonging to different classes. The models have been semi-automatically aligned to a consistent orientation using a hierarchical approach based on [25]. We mainly concentrated on car models, but we also trained a network on chairs to show generality of our approach and to allow a comparison to related methods. We used 7039 car models and 6742 chair models.

3D models were rendered using our self-developed real-time rendering framework based on the Panda3D rendering engine<sup>1</sup>. This allowed us to generate training images on the fly, without the need to store huge amounts of training data on the hard drive. We randomly sampled azimuth angles in the range from  $0^\circ$  to  $360^\circ$ , elevation angles in the range from  $-10^\circ$  to  $40^\circ$ , and the distance to the object from 1.7 to 2.3 units, with a car length being approximately equal to 3 units.

We took special care to ensure the realism of the renderings, since we would like the network to generalize to real input images. As in Su et al. [26], we randomly sampled the number of light sources from 2 to 4, each with random intensity and at random location. When overlaying the rendering on top of the background, we performed alpha compositioning to avoid sharp transition. It

<sup>1</sup> <https://www.panda3d.org>

was implemented by smoothing the segmentation mask with a Gaussian filter with the standard deviation randomly sampled between 1 and 1.3. Additionally, we smoothed the car image with a Gaussian filter with the standard deviation randomly sampled between 0.2 and 0.6.

Since we used a large amount of models, some of which may happen to be similar, simple random train-test splitting does not enable a reliable evaluation. To mitigate this problem we clustered objects according to their similarity and then took some of the clusters as the test set. We are mostly interested in splitting the objects according to their shape, so we used the distance between the 2D HOG descriptors [9] of the corresponding images as similarity measure. To make this measure more robust, we considered three different viewpoints for each object and used the sum of three distances as the final distance measure. After constructing a matrix of pairwise distances, we clustered the models using agglomerative clustering with average linkage.

For cars we selected a single cluster consisting of 127 models as the test set. Models from this group we refer to as 'normal test cars'. In addition, we picked 20 more models from the training set that have the highest distance from 'normal cars' and added them to the test set. Those are referred to as 'difficult test cars'. Example models from the test set and their corresponding nearest neighbors from the training set are shown in Figure 3. For chairs we picked three clusters as the test set comprising a total of 136 models.

## 4 Experimental evaluation

*Network training details* We used Tensorflow [27] for training the networks. The objective was optimized using the Adam method [28] with  $\beta_1 = 0.9$  and  $\beta_2 = 0.999$ . We initialized the weights of our network by sampling a Gaussian with corrected variance as described in [29]. The learning rate was equal to 0.0001.

We did not perform data augmentation, as we observed that it does not result in better generalization but leads to slower convergence. It seems there is already enough variation in the training data.


### 4.1 Unseen view prediction

We trained the networks to generate previously unseen views of objects from a single input image, therefore this is the first task we test on. Exemplary results for cars are shown in Figure 4. The network predicts the desired view for both normal and difficult (top right) cars, without (top row) and with (bottom row) background. The shape and the color of the car are always correctly estimated. Predictions for the difficult car are more blurry, since this car is dissimilar from models the network has been trained on.

Compared to the ground truth, the predictions are slightly blurry and lack some details. Apart from the fact that the problem is heavily ill-posed, this is



**Fig. 4.** Predictions of the network (top row for each model) and the corresponding ground truth images (bottom row for each model). The input to the network is in the leftmost column for each model. The top right model is a "difficult" car.



**Fig. 5.** Depth map predictions (**top row**) and the corresponding ground truth (**bottom row**). The network correctly estimates the shape.

likely to be a consequence of using squared Euclidean error as the loss function: if the network is uncertain about the prediction, it averages over potential images, resulting in blur. This could be mitigated for example by adversarial training, first proposed by Goodfellow et al. [30]. We experimented in this direction (see supplementary material for details), and indeed the images become slightly sharper, but at the cost of introduced noise, artifacts and very sensitive training. Therefore, we stick with the squared Euclidean loss.

**Comparison with a nearest neighbor baseline** We compare the network with a simple nearest neighbor (NN) baseline approach. We maximally simplify the task for the baseline approach: unlike the network it knows the input image viewpoint and there is no background. Given an input image with known viewpoint, the baseline searches the training set for the model which looks most similar from this viewpoint according to some metric. The prediction is simply the rendering of this model from the desired viewpoint. We tried three different metrics for the NN search: Euclidean distance in RGB space, Euclidean distance in HOG space, and a weighted combination of these.

	Color		Depth	
	Normal	Difficult	Normal	Difficult
NN HOG	0.028	0.039	0.0058	0.0225
NN HOG+RGB	0.020	0.036	0.0058	0.0221
NN RGB	0.018	0.034	0.0064	0.0265
Network	<b>0.013</b>	<b>0.028</b>	<b>0.0057</b>	<b>0.0207</b>

**Table 1.** Average error of predicted unseen views with our network and with the nearest neighbor baseline.

Table 1 reports average errors between the ground truth images and the predictions generated either with the baseline method or with our network. The error measure is Euclidean distance between the pixel values, averaged over the number of pixels in the image, the number of input and output viewpoints, the number of models and the maximum per pixel distance (443.4 for RGB and 65535 for depth). We separately show results for normal and difficult cars.


The network outperforms the baselines on both tasks, even though it is not given the input viewpoint. NN search can yield cars that look alike from the input view but may be very different when viewed from another angle. The network, in contrast, learns to find subtle cues which help to infer the 3D model. Another clear disadvantage of the NN search is that it can only return what is in the dataset, whereas the network can recombine the information of the training set to create new images.

**Comparison with existing work** We compared our results to several existing deep learning approaches that generate unseen views of images.

Except for a comparison to Dosovitskiy et al. [12], for which code was available, all comparisons are only on a qualitative basis. There are two reasons: first, there was no code to run other existing methods. Second, it is unclear which quantitative measure would be best to judge the quality of generated images. The best quantitative experiment would be a study with human observers, who have to assess which images look better. Since the differences in the quality of the results is mostly so obvious that quantitative numbers would not provide additional information, the lack of code is not a problem.

In order to compare with the Inverse Graphics Network (IGN) of Kulkarni et al. [8] we selected from our test set chair models similar to those Kulkarni et al. used for testing and showed in their paper. We also used the same input viewpoint. The results are shown in Figure 6. In all cases our network generates much more accurate predictions. Unlike IGN, it always predicts the correct view and generates visually realistic and detailed images. It captures fine details like bent legs in the top example or armrests in the second example.

We also compared to Dosovitskiy et al. [12]. This approach allows the prediction of all views of a chair model given only a single view during training. However, 1) it requires several days of training to be able to predict unseen



**Fig. 6.** Our results (black background) compared with those from IGN [8] (white background) on similar chair models. The leftmost image in each row is the input to the network. In all cases our results are much better.




**Fig. 7.** Comparison of our approach (top row for each model) with novel view prediction results from Dosovitskiy et al. [12] (bottom row for each model). The estimates of our network are more accurate and consistent, and it does not require re-training for each new model.

views of a new chair and 2) it is not explicitly trained to predict these unseen views, so there is no guarantee that the predictions would be good. We used the code provided by the authors to perform comparisons shown in Figure 7. For each model the top row shows our predictions and the bottom row those from Dosovitskiy et al. While in simple cases the results look qualitatively similar (top example), our approach better models the details of the chair style (chair legs in the bottom example). This is supported by the numbers: the average error of the images predicted by our network is 0.0208 on the chairs dataset, whereas the network of Dosovitskiy et al. has an average error of 0.0308. The error measure is the same as in the baseline comparison.

Finally, we qualitatively compared our results with the recent work of Yang et al. [14]. Here we show the results on cars, which we found to be more challenging than chairs. Figure 8 shows predictions by Yang et al. (top row for each model) and our work (bottom row for each model). For both models the leftmost column shows the input image. We picked the cars from our dataset that most resemble the cars depicted in Yang et al. [14]. Since images generated by their method are  $64 \times 64$  pixels, we downsampled our results to this resolution to make the visual comparison fair. Our predictions look much more realistic and significantly less blurred. The method of Yang et al. occasionally averages several viewpoints (for example, the third column from the left for the top model in Figure 8), while our method always generates sharp images as seen from the desired viewpoint.

**Natural input images** To verify the generalization properties of our network, we fed it with images of real cars downsampled to the size of  $128 \times 128$  pixels. The results are shown in Figure 9. We do not have ground truth for these images so only the output of the network is shown. The quality of the predictions is slightly worse than for the (synthetic) test cars. The reasons may be complicated reflections and camera models different from ones we used for rendering. Still, the network estimates the shape and the color well.



**Fig. 8.** Predictions from Yang et al. [14] (top row for each model) compared to our predictions (bottom row for each model) on similar car models. The leftmost image in each row is the input to the network. Our network generates more realistic and less blurred images.




**Fig. 9.** Network predictions for natural input images. The net correctly estimates the shape and color.

We observed that the realistic rendering procedure we implemented is important for generalization to real images. We show in the supplementary material that simpler rendering leads to complete failure on real data.

We emphasize that this is the first time neural networks are shown to be able to infer 3D representations of objects from real images. Interesting avenues of future research include deeper study of the network’s performance on real images, as well as joint training on synthetic and real data and applications of transfer learning techniques. However, these are beyond the scope of this paper.

## 4.2 3D model prediction

We verified to which extent the predicted depth maps can be used to reconstruct full 3D surfaces. Figure 5 shows two exemplary depth maps generated by our network together with the corresponding ground truth. The overall quality is



**Fig. 10.** 3D model reconstructions of a "normal" (left) and a "difficult" (right) car.

similar to that of predicted images: the shape is captured correctly while some fine details are missing.


In Figure 10 we show 3D models obtained by fusing 6 predicted depth maps ( $\theta^{el} = 20^\circ$ ,  $\theta^{az} = \{0^\circ, 60^\circ, 120^\circ, 180^\circ, 240^\circ, 300^\circ\}$ ). Already the raw point clouds represent the shape and the color well, even for the "difficult" model (right). Dense depth map fusion removes the noise and a smooth surfaces, yet also destroys some more details due to the regularizer involved. For more results on 3D models we refer to the video in the supplemental material.

### 4.3 Analysis of the network

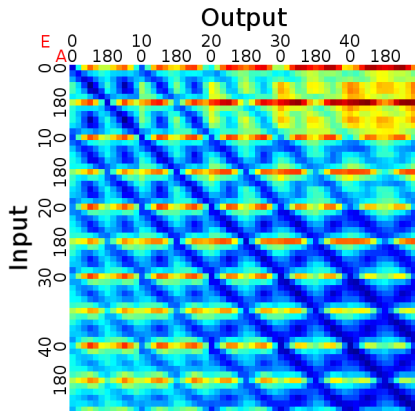
**Viewpoint dependency** Since the prediction task is ambiguous, the quality of predictions depends on how informative the input image is with regard to the desired output view. For our network we can observe this tendency, as shown in Figure 11. If the input viewpoint reveals much about the shape of the car, such as the side-view input, the generated images match the ground truth quite well. In case of less informative input, such as the front-view input, the network has to do more guesswork and resorts to predicting the most probable answer. However, even if the input image is weakly informative, all the predicted views correspond to a consistent 3D shape, indicating that the network first extracts a 3D representation from the image and then renders it from different viewpoints.

In Figure 12 we quantify the prediction quality depending on the input and output views. The matrix shows the Euclidean distance between the generated and ground truth images for different input (y-axis) and output views (x-axis) averaged over the whole test set. Each column is normalized by its sum to compensate for different numbers of object pixels in different views. Several interesting patterns can be observed. The prediction task gets harder if the input view is very different from the output view, especially if the input elevation is small: top right and bottom left corners of the matrix are higher than the rest. Local patterns show that for each elevation it is easier to predict images with the same





**Fig. 11.** The more informative the input view is, the better the network can estimate the ground truth image. For uninformative inputs it simply invents some model which is still internally consistent.



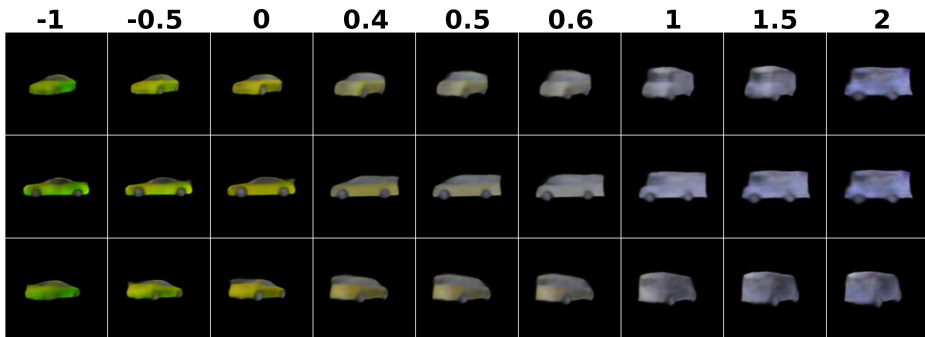
**Fig. 12.** Distance from ground truth for different input and output views. Shown are all combinations of elevation (E) and azimuth (A) angles with a  $30^\circ$  step. It is harder to predict views that are significantly different from the input.

or similar azimuth angles. Diagonal blue stripes show that it is easier to predict similar or symmetric views.


**Object interpolation** The hidden object representation extracted by the network is not directly interpretable. One way to understand it is to modify it and see how this affects the generated image. In the experiment shown in Figure 13, we encoded two extremely different models (a car and a bus) into feature vectors  $f_{car}$  and  $f_{bus}$ , linearly interpolated between these  $f_{int} = \alpha f_{car} + (1 - \alpha) f_{bus}$ , and decoded the resulting feature vectors. We also tried extrapolation, that is,  $\alpha < 0$  and  $\alpha > 1$ .

The first and most important observation is that all generated views form consistent shapes, which strongly indicates that the interpolation modifies the 3D representation, which is then rendered from the desired viewpoint. Second, extrapolation also works well, exaggerating the ‘carness’ or the ‘busness’ of the models. Third, we observed that the morphing is not uniform: there is not much happening for  $\alpha$  values close to 0 and 1, most of the changes can be seen when  $\alpha$  is around 0.5.

**Internal representation** In order to study the properties of the internal representation of the network, we ran the t-SNE embedding algorithm [31] on the 1024-dimensional vectors computed for a random subset of models from the training set with fixed viewpoint. t-SNE projects high-dimensional samples to a 2D space such that similar samples are placed close to one another. The results



**Fig. 13.** Morphing a car into a bus by interpolating between the feature representations of those two models. All the intermediate models are consistent.



**Fig. 14.** t-SNE embedding in latent 1024-dimensional space. Cars are grouped according to their shape and color.

of this experiment are shown in Figure 14. Both shape and color are important, but shape seems to have more weight: similar shapes end up close in the 2D space and are sorted by color within the resulting groups.

In Section 3 of the supplementary material we also show that different input views of the same object lead to very similar intermediate representations.

## 5 Conclusions

We have presented a feed-forward network that learns implicit 3D representations when being trained on the task to generate new views from a single input image. Apart from rendering any desired view of an object, the network allows us to also generate a point cloud and a surface mesh. Although the network was trained only on synthetic data, it can also take natural images as input. Clearly, natural images are harder for the network since the training data does not yet fully model all variations that appear in such images. In future work we will investigate ways to improve the training set either by more realistic renderings or by ways to mix in real images.

## 6 Acknowledgments

We acknowledge funding by the ERC Starting Grant VideoLearn (279401). We would like to thank Nikolaus Mayer and Benjamin Ummenhofer for their comments.

## References

1. Savva, M., Chang, A.X., Hanrahan, P.: Semantically-Enriched 3D Models for Common-sense Knowledge. CVPR 2015 Workshop on Functionality, Physics, Intentionality and Causality (2015)
2. Memisevic, R., Hinton, G.: Unsupervised learning of image transformations. In: CVPR. (2007)
3. Michalski, V., Memisevic, R., Konda, K.R.: Modeling deep temporal dependencies with recurrent grammar cells”. In: NIPS. (2014) 1925–1933
4. Reed, S., Sohn, K., Zhang, Y., Lee, H.: Learning to disentangle factors of variation with manifold interaction. In: ICML. (2014)
5. Lee, H., Grosse, R., Ranganath, R., Ng, A.Y.: Convolutional deep belief networks for scalable unsupervised learning of hierarchical representations. In: ICML. (2009) 609–616
6. Hinton, G.E., Krizhevsky, A., Wang, S.D.: Transforming auto-encoders. In: ICANN. (2011) 44–51
7. Zhu, Z., Luo, P., Wang, X., Tang, X.: Multi-view perceptron: a deep model for learning face identity and view representations. In: NIPS. (2014) 217–225
8. Kulkarni, T.D., Whitney, W.F., Kohli, P., Tenenbaum, J.: Deep convolutional inverse graphics network. In: Advances in Neural Information Processing Systems (NIPS) 28. (2015) 2539–2547
9. Dalal, N., Triggs, B.: Histograms of oriented gradients for human detection. In: International Conference on Computer Vision & Pattern Recognition. Volume 2. (June 2005) 886–893
10. Chen, C.Y., Grauman, K.: Inferring unseen views of people. In: CVPR. (2014)
11. Su, H., Wang, F., Yi, L., Guibas, L.J.: 3d-assisted image feature synthesis for novel views of an object. In: ICCV. (2015)
12. A.Dosovitskiy, J.T.Springenberg, T.Brox: Learning to generate chairs with convolutional neural networks. In: IEEE International Conference on Computer Vision and Pattern Recognition (CVPR). (2015)
13. Yang, J., Reed, S.E., Yang, M.H., Lee, H.: Weakly-supervised disentangling with recurrent transformations for 3d view synthesis. In Cortes, C., Lawrence, N.D., Lee, D.D., Sugiyama, M., Garnett, R., eds.: Advances in Neural Information Processing Systems 28. Curran Associates, Inc. (2015) 1099–1107
14. Yang, J., Reed, S., Yang, M., Lee, H.: Weakly-supervised disentangling with recurrent transformations for 3d view synthesis. arXiv:1601.00706 (2016)
15. Barron, J.T., Malik, J.: Shape, illumination, and reflectance from shading. TPAMI (2015)
16. Blanz, V., Vetter, T.: Face recognition based on fitting a 3D morphable model. TPAMI 25(9) (September 2003)
17. Liu, F., Zeng, D., Li, J., Zhao, Q.: Cascaded regressor based 3d face reconstruction from a single arbitrary view image. arXiv:1509.06161 (2015)

18. Kar, A., Tulsiani, S., Carreira, J., Malik, J.: Category-specific object reconstruction from a single image. In: CVPR. (2015) 1966–1974
19. Vicente, S., Carreira, J., de Agapito, L., Batista, J.: Reconstructing PASCAL VOC. In: CVPR. (2014) 41–48
20. Aubry, M., Maturana, D., Efros, A., Russell, B., Sivic, J.: Seeing 3d chairs: exemplar part-based 2d-3d alignment using a large dataset of cad models. In: CVPR. (2014)
21. Huang, Q., Wang, H., Koltun, V.: Single-view reconstruction via joint analysis of image and shape collections. *ACM Trans. Graph.* **34**(4) (2015) 87
22. Eigen, D., Puhrsch, C., Fergus, R.: Depth map prediction from a single image using a multi-scale deep network. In: NIPS. (2014)
23. Wu, Z., Song, S., Khosla, A., Yu, F., Zhang, L., Tang, X., Xiao, J.: 3d shapenets: A deep representation for volumetric shapes. In: CVPR. (2015) 1912–1920
24. Pock, T., Zebedin, L., Bischof, H.: Tgv-fusion. In: *Rainbow of Computer Science*. Volume 6570 of *Lecture Notes in Computer Science*. (2011) 245–258
25. Huang, Q.X., Su, H., Guibas, L.: Fine-grained semi-supervised labeling of large shape collections. *ACM Trans. Graph.* **32**(6) (November 2013) 190:1–190:10
26. Su, H., Qi, C.R., Li, Y., Guibas, L.J.: Render for cnn: Viewpoint estimation in images using cnns trained with rendered 3d model views. In: *The IEEE International Conference on Computer Vision (ICCV)*. (December 2015)
27. Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G.S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., Zheng, X.: TensorFlow: Large-scale machine learning on heterogeneous systems (2015) Software available from tensorflow.org.
28. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: *ICLR*. (2015)
29. He, K., Zhang, X., Ren, S., Sun, J.: Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In: *2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, December 7-13, 2015*. (2015) 1026–1034
30. Goodfellow, I.J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A.C., Bengio, Y.: Generative adversarial nets. In: *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*. (2014) 2672–2680
31. van der Maaten, L., Hinton, G.: Visualizing high-dimensional data using t-sne. *Journal of Machine Learning Research* **9** (2008) 2579–2605