

Deep Textured 3D Reconstruction of Human Bodies

Abbhinav Venkat
abbhinav.venkat@research.iiit.ac.in

Sai Sagar Jinka
jinka.sagar@reseach.iiit.ac.in

Avinash Sharma
asharma@iiit.ac.in

Center for Visual Information
Technology (CVIT),
Kohli Center for Intelligent Systems
(KCIS),
IIIT Hyderabad,
India

Abstract

Recovering textured 3D models of non-rigid human body shapes is challenging due to self-occlusions caused by complex body poses and shapes, clothing obstructions, lack of surface texture, background clutter, sparse set of cameras with non-overlapping fields of view, etc. Further, a calibration-free environment adds additional complexity to both - reconstruction and texture recovery. In this paper, we propose a deep learning based solution for textured 3D reconstruction of human body shapes from a single view RGB image. This is achieved by first recovering the volumetric grid of the non-rigid human body given a single view RGB image followed by orthographic texture view synthesis using the respective depth projection of the reconstructed (volumetric) shape and input RGB image. We propose to co-learn the depth information readily available with affordable RGBD sensors (e.g., Kinect) while showing multiple views of the same object during the training phase. We show superior reconstruction performance in terms of quantitative and qualitative results, on both, publicly available datasets (by simulating the depth channel with virtual Kinect) as well as real RGBD data collected with our calibrated multi Kinect setup.

1 Introduction

Recovering the textured 3D model of non-rigid human shapes from images is of high practical importance in the entertainment industry, e-commerce, health care (physiotherapy), mobile based AR/VR platforms, etc. This is a challenging task as the object geometry of non-rigid human shapes evolve over time, yielding a large space of complex body poses as well as shape variations. In addition to this, there are several other challenges such as self-occlusions by body parts, obstructions due to free form clothing, background clutter (in a non-studio setup), sparse set of cameras with non-overlapping fields of views, sensor noise, etc. Figure 1 illustrates some of these challenges. Model based reconstruction techniques attempt to overcome some of these limitations, however, at the cost of loss of accurate geometrical information over the shape surface and are typically applicable only for tight clothing scenarios [2, 3, 4].

Traditionally, calibrated multi-camera setups have been employed to recover textured 3D models through triangulation or voxel carving [5, 6]. However, these techniques yield




Figure 1: Challenges in non-rigid reconstruction. (a) Complex poses (b) Clothing obstructions (c) Shape variations (d) Background clutter.

reconstructions with severe topological noise [12]. Recent attempts replace/augment the capture setup with high resolution depth sensors [1, 24], thus making the setup more accurate but less affordable. Nevertheless, the fundamental limitation of these techniques is the requirement of a calibrated multi-camera/sensors that restricts their applicability to studio environments. We aim at achieving 3D reconstruction of textured human body models using single/multiple affordable RGB/D sensor(s) in a calibration-free setup.

1.1 Related Work

Recent advancement in deep networks has enabled learning a class specific 3D structure of a set of objects (e.g. cars, chairs, rooms) using large scale datasets of synthetic 3D models for training the deep networks [6, 13, 19, 20, 21]. ShapeNet[20] proposed a deep network representation with a convolutional deep belief network to give a probabilistic representation of the voxel grid. Along similar lines, 3D Generative Adversarial Networks (GAN’s) were proposed to learn a probabilistic latent space of rigid objects (such as chairs, tables) in [19]. [21] proposed an encoder-decoder network that utilizes observations from the 2D space, and without supervision from 3D, performs reconstruction for a few classes of objects. This relationship between 2D and 3D is further exploited in [15] where they define a loss function in terms of ray consistency and train for single view reconstruction. However, these methods have been employed only for rigid object reconstruction.

In regard to non-rigid reconstruction, [14] proposed to directly achieve the surface reconstruction by learning a mapping between 3D shapes and the geometry image representation (introduced in [3]). One of the key limitations of their method is that it is only suitable for genus-0 meshes. This constraint is frequently violated in real-world human body shapes due to topological noise [12] induced by complex poses, clothing, etc. Another very recent work proposed in [18] uses multi-view silhouette images to obtain reconstructions of free form blob-like objects/sculptures. However, the use of silhouettes limits the application to scenarios where background subtraction is assumed to be trivial. All these initial efforts are focused on textureless 3D reconstruction and do not seem to directly extend-able to non-rigid human body shapes. DynamicFusion [11] proposed the first dense SLAM system capable of reconstructing non-rigidly deforming scenes in real-time, by fusing together RGBD scans captured from commodity sensors. Similar to them, we shall be training our network with data from RGB/D sensors, however, unlike them, our pipeline shall be capable of high quality reconstruction from a single image by exploiting human body symmetry. Just as is true in their case, the reconstruction quality improves with providing more view information.

In regard to texture generation, recent work on multi-view synthesis [22, 23] propose to generate an image from an alternate view-point given an input image. Recent work proposed in [8] attempts texture recovery for category specific 3D models. However, the texture map is predicted only for the mean shape in the UV space and shown on reconstruction of birds. Nevertheless, this is not easily generalizable to human body shapes as accurate warping is non-trivial due to a large space of shape and pose variations associated with the human body. Very recently, a model based learning paradigm outlined in [10] proposed to generate human body models by estimating their deviation from the nearest SMPL model. Nevertheless, to recover the texture, the person is assumed to be seen from all the sides in a video sequence. Additionally such model based methods fail to deal with large geometric deformations induced by free-form clothing (e.g., rob or long skirt) scenarios.

1.2 Our Contribution

In this paper, we propose a deep learning based solution for textured 3D reconstruction of human body shapes given an input RGB image, in a calibration-free environment. Given a single view RGB image, both reconstruction and texture generation are ill-posed problems. Thus, we proposed to co-learn the depth cues (using depth images obtained from affordable sensors like Kinect) with RGB images while training the network. This helps the network learn the space of complex body poses, which otherwise is difficult with just 2D content in RGB images. Although we propose to learn the reconstruction network with multi-view RGB and depth images (shown one at a time during training), co-learning them with shared filters enabled us to recover 3D volumetric shapes using just a single RGB image at test time. Apart from the challenge of non-rigid poses, the depth information also helps addressing the challenges caused by cluttered background, shape variations and free form clothing. Our texture recovery network uses a variational auto-encoder to generate orthographic texture images of reconstructed body models that are subsequently backprojected to recover a texture 3D mesh model. We show quantitative and qualitative results on four publicly available datasets (by simulating the RGB and D whenever unavailable) as well as real RGBD data collected with calibrated multi Kinect setup. The key contributions of this work are:

- First, we introduce a novel deep learning pipeline to obtain textured 3D models of non-rigid human body shapes from a single image. To the best of our knowledge, obtaining the reconstruction of non-rigid shapes in a volumetric form (whose advantages we demonstrate) has not yet been attempted in literature. Further, this would be an initial effort in the direction of single view non-rigid reconstruction and texture recovery in an end-to-end manner.
- Second, we demonstrate the importance of depth cues (*used only at train time*) for the task of non-rigid reconstruction. This is achieved by our novel training methodology of alternating RGB and D in order to capture the large space of pose and shape deformation.
- Third, we show that our model can partially handle non-rigid deformations induced by free form clothing, as we do not impose any model constraint while training the volumetric reconstruction network.
- Fourth, we proposed to use depth cues for texture recovery in the variational auto-encoder setup. This is the first attempt to do so in texture synthesis literature.




Figure 2: Proposed end-to-end deep learning pipeline for reconstructing textured non-rigid 3D human body models. Using a single view perspective RGB image (a), we perform a voxelized 3D reconstruction (c) using the reconstruction network (b). Then, to add texture to the generated 3D model, we first convert the voxels to a mesh representation using Poisson’s surface reconstruction algorithm [9] and capture its four orthographic depth maps (d). These are fed as an input to the texture recovery network (f), along with the perspective RGB views used for reconstruction (a). The texture recovery network produces orthographic RGB images (g), that are back-projected onto the reconstructed model, to obtain the textured 3D model (h).

- Finally, we collected a real dataset (that shall be publicly released) of textured 3D human body models and their corresponding multi-view RGBD, that can be used in solving a variety of other problems such as human tracking, segmentation etc.

2 Our Method

A naive method of obtaining a textured 3D model would be to predict a colored occupancy grid. Given that the size of the volumetric output is 128^3 and that for each such voxel, we would have to predict three values - one each for R, G and B (each in the range of 0 – 255), it would be extremely computationally expensive to follow this method. Hence, we break down our end-to-end pipeline into two stages - one for 3D reconstruction and another for texture synthesis.

Figure 2 shows the test time flow of the proposed end-to-end pipeline for reconstruction of textured non-rigid 3D shapes. A detailed outline of each stage is given below.

2.1 Non-rigid Reconstruction

For non-rigid reconstruction, we propose to use an encoder-decoder network with two methods for combining multi-view information - i) a 3D GRU (like in [8]) and ii) max pooling of the CNN feature vectors of the views (like in [18]). The above two settings are considered to show that irrespective of the method of combining multi-view information for non-rigid human body reconstruction, depth information immensely help to capture the vast range of

complex pose and shape variations while enhancing the reconstruction accuracy. Moreover, this provides the added advantage of obtaining 3D reconstruction from either only RGB or only D at test time, as discussed below. Further, by randomizing the number of views used at train time, we enable the network to reconstruct from a single image at test time. Intuitively, the network exploits the symmetry associated with non-rigid human shapes to reconstruct their 3D, even from a single view.

Network Architecture.

- a) *Encoder* - ResNet-18 is used here that takes images of size 255x255 in one of the input modes (see below) and produces a 1024 dimensional feature vector. Each view produces one such feature vector, which is combined in the multi-view module. We use Leaky ReLU as the activation function in both the encoder and decoder.
- b) *Multi-view Module* - Multi-view information is combined using either a max-pooling of the 1D CNN feature vectors, or using a 3D-GRU. The outputs are resized to 4^3 and fed to the decoder
- c) *Decoder* - Deconvolutional ResNet-18 is used here that up-samples the output of the multi-view module to 128^3

Input Modes. In order to capture the large space of complex pose and shape deformations of humans, we experiment with four input modes:

- a) *RGB* - This setup is commonly used in rigid body reconstructions [6, 14]. However, we qualitatively and quantitatively show in Section 3.2 that this setup is inadequate for reconstructing non-rigid shapes.
- b) *D* - The premise behind this mode is that depth-maps give us information about the geometry of the object, which as seen in Section 3.2, help in significantly enhancing the reconstruction quality.
- c) *RGBD* - In order to exploit both the depth and color channels, we augment RGB with D in a 4 channel input setup.
- d) *RGB/D* - Lastly, we propose a unique training methodology that gives us superior performance than the above 3 modes. Here, at train time, each mesh is reconstructed from either only multi-view RGB or multi-view depth. Thus, while we train with depth information, we can test with only RGB as well, which is a major advantage. Intuitively, this strategy is equivalent to sharing weights between the RGB and D spaces, in order to exploit the coherence between RGB and D; thus combining the advantages from both the spaces.

Loss Function. We use Voxel-wise Cross Entropy to train the reconstruction models. It is the sum of the cross-entropies for each pair of predicted and target voxel values. Let ' p ' be the predicted value at voxel (i, j, k) and ' y ' the corresponding expected value. Then, the loss is defined as :

$$L(p, y) = \sum_{i,j,k} y_{(i,j,k)} \log(p_{(i,j,k)}) + (1 - y_{(i,j,k)}) \log(1 - p_{(i,j,k)}) \quad (1)$$

2.2 Texture Recovery

We obtain a set of four orthographic depth maps (excluding top and bottom views) of the 3D human mesh M generated by the reconstruction network and represent them as $D = \{\pi_{D,i}^o\}$. Here $\pi_{D,i}^o$ corresponds to the orthographic (o) depth projection (D) of M to the i^{th} face of

the cube. From the set of perspective RGB images used for reconstruction $P = \pi_{RGB,j}^P$, we randomly choose an image as an input to the texture recovery network. Given each of the orthographic depth images and a randomly chosen perspective image, texture recovery aims to estimate the color of each pixel in the depth map.

In order to achieve this, we propose a simple texture recovery approach that detaches the requirement of a calibrated setup by employing a variational auto-encoder (VAE). Another generative model, generative adversarial networks (GAN) can be employed, however, we leverage upon VAEs since it is better at understanding the global shape of the object while GANs are traditionally used for recovering finer details. The VAE is trained by maximizing the log likelihood $\log(p_\theta(\pi_{RGB,i}^O | \pi_{D,i}^O, \pi_{RGB,j}^P))$ which is equivalent to minimizing the variational lower bound, as indicated by the equation below. The VAE is trained to learn the distribution $p(\pi_{RGB,i}^O | \pi_{D,i}^O, \pi_{RGB,j}^P)$ which models the color of the provided orthographic depth maps.

$$V(\pi_{RGB,i}^O, \pi_{D,i}^O, \pi_{RGB,j}^P; \theta, \phi) = -KL(q_\phi(z | \pi_{D,i}^O, \pi_{RGB,j}^P) || p_\theta(z)) + E_{q_\phi}[\log p_\theta(\pi_{RGB,i}^O | \pi_{D,i}^O, \pi_{RGB,j}^P, z)] \quad (2)$$

The data is modeled by normalizing out z from the joint probability distribution $p(\pi_{RGB,i}^O, z)$. Further, $p(z)$ is inferred as $p_\theta(z | \pi_{D,i}^O, \pi_{RGB,j}^P)$ and is modeled by the encoder, parameterized by θ . Variational distribution q_ϕ is introduced to approximate the unknown true posterior p_θ . For further details, refer to the formulation of Variational Auto-encoders [10]. The second term in equation 2 is the log likelihood of samples, which is the L2 reconstruction loss that encourages consistency between the encoder and decoder. The KL divergence reduces the distance between the variational distribution and prior distribution. $q_\phi(z|x)$ is associated with a prior distribution over the latent variables for which we use a multivariate Gaussian with unit variance $\mathcal{N}(0, I)$. To the best of our knowledge, this is the first attempt to use depth cues for novel view synthesis.

Network Architecture. As shown in Figure 2, the encoder consists of two symmetric branches, one each, for orthographic depth, and perspective RGB image as input. However, the weights are not shared between the two encoders. These networks consist of convolutional layers which have 64,128,256,256,512 and 1024 channels with filters of size 5,5,3,3,3 and 4 respectively followed by a fully connected layer with 1024 neurons. The representations from the two branches are combined to form a 1024 dimensional latent variable. The decoder network consists of a fully connected layer with $256 \times 8 \times 8$ neurons. Deconvolution layers are then followed with 2×2 upsampling and have 256, 256, 256, 128, 64 and 3 channels. The filter sizes are 5×5 for all layers. ReLU is used as the activation function and the output size is set to 64×64 .

3 Experiments & Results

3.1 Datasets

MPI Datasets: We use two datasets from MPI. First, we use the parametric SMPL model [1, 11] to generate 10 mesh sequences, each containing 300 frames, i.e., 3000 meshes, consisting of an equal number of male and female models. Additionally, we use FAUST’s data [12]




Figure 3: 3D Shapes obtained with our reconstruction network compared to ground truth models obtained with MVG.

Dataset	Multi-View	RGB (Baseline)	D	RGBD	RGB/D (Ours)
MPI-SMPL [16]	3D-GRU	0.6903	0.7709	0.7541	0.8040
	Max Pool	0.7144	0.7633	0.7550	0.7816
MIT [17]	3D-GRU	0.0103	0.7403	-	0.7547
	Max Pool	0.0081	0.7205	-	0.7480
MPI-FAUST [8]	3D-GRU	0.8113	0.8629	0.8356	0.8644
	Max Pool	0.8150	0.8661	0.8366	0.8521
OUR DATA	3D-GRU	0.6816	0.7963	0.8114	0.8241
	Max Pool	0.6883	0.7844	0.8017	0.8066

Table 1: A comparison of IoU values tested using a single view on datasets [8, 16, 17], under the various input modes, when trained with two different view modules.

which consists of 300 high resolution human meshes. There are a total of 10 different subjects in 30 different poses. The 300 meshes come divided in 2 sequences, one having complete meshes and the other having broken/incomplete parts - the former used for training, and the latter for testing. In addition, we simulated a virtual Kinect setup to capture aligned RGBD. Since both these datasets had correspondences, a few meshes were manually textured, and this texture was transferred across each of the datasets.

MIT’s Articulated Mesh Animation [17]: This dataset consists of 5 mesh sequences (approx. 175 to 250 frames each). It provides RGB images from 8 views and the corresponding 3D meshes for each frame. The total number of meshes used from this dataset for training are 1,525.

Our Data: A total of 5 mesh sequences, each containing 200 to 300 frames with significant pose and shape variation were captured using a calibrated multi-Kinect setup. The colored point clouds were re-meshed using Poisson’s surface reconstruction algorithm, after pre-processing them for noise removal. The processed RGBD and corresponding textured models were used for training the pipeline.

3.2 Non-rigid Reconstruction

Network’s Training. We used Nvidia’s GTX 1080Ti, with 11GB of VRAM to train our models. A batch size of 5 with the ADAM optimizer having an initial learning rate of 10^{-4} , and a weight decay of 10^{-5} is used to get optimal performance. Further, a standard 80 : 20

split between the training and testing datasets is adhered to. In order to ensure that reconstruction is feasible from single as well as multiple views, we choose random number of views from available views for training a mesh in each iteration. Using this randomization in training, we are providing sufficient view information to the network so that it can learn the space of body pose and shape variations and hence able to achieve single view reconstruction at test time.

Evaluation Metric. The primary metric used to evaluate our performance is the Intersection over Union (IoU), which is a comparison between the area of overlap and the total area encompassing both the objects. Larger its value, the better the quality of reconstruction.

Let ' p ' be the predicted value at voxel (i, j, k) and ' y ' the corresponding expected value. ' I ' is an indicator function which gives a value of 1 if the expression it is evaluating is true, if not, it gives 0. ' t ' is an empirically decided threshold of 0.5 above which the cell is considered as filled.

$$IoU = \frac{\sum_{i,j,k} [I(p(i, j, k) > t)I(y(i, j, k))]}{\sum_{i,j,k} [I(p(i, j, k) > t) + I(y(i, j, k))]} \quad (3)$$



Figure 4: Clothing induced deformations captured by our proposed method on [17].

Baseline. As described in Section 1.1, there are several standard encoder-decoder networks that use single/multi-view RGB image(s) for voxelized reconstruction of rigid objects. Therefore, as baseline, we interpret this setting of using only RGB image(s) for both training and testing the reconstruction network. Further, the qualitative and quantitative results are compared and evaluated on ground truth data generated using traditional Multi-View Geometry (MVG) setups (from upto 22 views), on four datasets of varying complexity. For rendered sample outputs, please refer to the supplementary video.

Results & Discussion. Quantitative results (IOU metric) in Table 1 suggest that for a variety of datasets of varying complexity and irrespective of the method of combining multiple views, the depth information is very critical for accurate reconstruction of human models. It is interesting to notice that the difference in IoU values between RGB and RGB/D widens under two scenarios - a) when the dataset has very complicated poses (such as the handstand sequence in MIT) and b) when the background becomes more complicated. Figure 5 re-emphasizes this ineffectiveness of using RGB alone, in which the first and second rows show the reconstructions from MIT's handstand and our captured data, respectively. An intuition behind the working of this training paradigm is that the co-learning of the shared filter weights of the two modalities act as a regularization for one another, thus enhancing the information seen by the network.

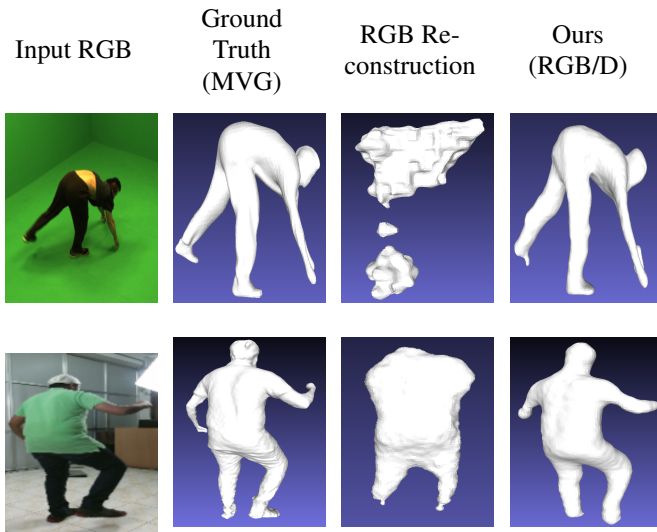


Figure 5: Qualitative comparison with baseline method.

Figure 3 shows the robustness of the learned model performing a vast range of actions. This robustness while reconstructing from a single image can be attributed to the network’s ability to exploit the symmetry associated with non-rigid human shapes. As a result of not imposing any body model constraint, we were able to partially handle non-rigid deformations induced by free form clothing as shown in Figure 4. While the current pipeline has not been trained to explicitly capture the temporal information available in sequences, results on MIT’s Samba dance sequence [17] shown in Figure 4 show great promise for reconstructing performance capture scenarios from a single image.

3.3 Texture Recovery

The textured model is obtained by synthesizing a texture orthographic view image corresponding to respective orthographic depth image taken from volumetric model predicted by reconstruction module.

Results & Discussion. Figure 6 shows the generated orthographic texture image and corresponding textured model obtained after their back-projection on reconstructed 3D model on samples from various datasets. For rendered textured meshes, please refer to the supplementary video. The variational autoencoder predicts low resolution images of size 64×64 . Attempting to predict higher resolution texture images using only this framework adds additional complexity to the network, and results in less accurate orthographic texture images. Further post-processing using super-resolution networks, usage of multi-view input images etc. can be used to improve the resolution of the synthesized texture images to perform high quality texture recovery.

Since we use just a single image to recovery the texture, the nature of the textures recovered shall be simple. Recovery of complex textures from only a single image is not a realistic expectation for any texture synthesis methodology. However, the task of single view texture

recovery has practical applications in performance capture scenarios, AR/VR platforms etc., and this is an initial effort to address that gap. Given a prior on the type of clothing or sequence that is going to be recovered (for example, if we would like to recover the texture of a particular dance form), the proposed texture recovery pipeline shall intuitively perform better (as it currently trained for multiple different scenarios).




Figure 6: Results on reconstruction and texture recovery on [4, 7] and on our real data.

4 Conclusion

We proposed a novel deep learning pipeline for reconstructing textured 3D models of non-rigid human body shapes using single view RGB images (during test time). This is a severely ill posed problem due to self-occlusions caused by complex body poses and shapes, clothing obstructions, lack of surface texture, background clutter, single view, etc. We show superior textured reconstruction performance using the proposed method in terms of quantitative and qualitative results on both publicly available datasets (by simulating the depth channel with virtual Kinect) as well as real RGBD data collected with a calibrated multi Kinect setup. As a part of future work, it will be practical to extend this to recovering high resolution textured models and also exploiting the temporal consistency for the task of reconstruction, for the case of continuous mesh sequences.

References

- [1] Thiemo Alldieck, Marcus A. Magnor, Weipeng Xu, Christian Theobalt, and Gerard Pons-Moll. Video based reconstruction of 3d people models. *CoRR*, abs/1803.04758, 2018. URL <http://arxiv.org/abs/1803.04758>.
- [2] Dragomir Anguelov, Praveen Srinivasan, Daphne Koller, Sebastian Thrun, Jim Rodgers, and James Davis. Scape: shape completion and animation of people. *ACM Transaction on Graphics*, 24:408–416, 2005.
- [3] Federica Bogo, Javier Romero, Matthew Loper, and Michael J. Black. FAUST: Dataset and evaluation for 3D mesh registration. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, Piscataway, NJ, USA, June 2014. IEEE.
- [4] Federica Bogo, Angjoo Kanazawa, Christoph Lassner, Peter Gehler, Javier Romero, and Michael J. Black. Keep it SMPL: Automatic estimation of 3D human pose and shape from a single image. In *European Conference on Computer Vision (ECCV)*, 2016.
- [5] Federica Bogo, Javier Romero, Gerard Pons-Moll, and Michael J. Black. Dynamic FAUST: Registering human bodies in motion. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [6] Christopher B. Choy, Danfei Xu, JunYoung Gwak, Kevin Chen, and Silvio Savarese. 3d-r2n2: A unified approach for single and multi-view 3d object reconstruction. In Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling, editors, *European Conference on Computer Vision (ECCV)*, pages 628–644, 2016.
- [7] Mingsong Dou, Sameh Khamis, Yury Degtyarev, Philip Davidson, Sean Ryan Fanello, Adarsh Kowdle, Sergio Orts Escolano, Christoph Rhemann, David Kim, Jonathan Taylor, Pushmeet Kohli, Vladimir Tankovich, and Shahram Izadi. Fusion4d: Real-time performance capture of challenging scenes. *ACM Transaction on Graphics*, 35(4): 114:1–114:13, July 2016.
- [8] Angjoo Kanazawa, Shubham Tulsiani, Alexei A. Efros, and Jitendra Malik. Learning category-specific mesh reconstruction from image collections. *arXiv preprint arXiv:1803.07549*, 2018.
- [9] Michael Kazhdan and Hugues Hoppe. Screened poisson surface reconstruction. *ACM Transactions on Graphics (TOG)*, 32(3):29, 2013.
- [10] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [11] Richard A Newcombe, Dieter Fox, and Steven M Seitz. Dynamicfusion: Reconstruction and tracking of non-rigid scenes in real-time. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 343–352, 2015.
- [12] Avinash Sharma, Radu Horaud, Jan Cech, and Edmond Boyer. Topologically-robust 3d shape matching based on diffusion geometry and seed growing. In *Computer Vision and Pattern Recognition (CVPR)*, pages 2481–2488, 2011.

- [13] Ayan Sinha, Jing Bai, and Karthik Ramani. Deep learning 3d shape surfaces using geometry images. In Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling, editors, *European Conference on Computer Vision (ECCV)*, pages 223–240, 2016.
- [14] Ayan Sinha, Asim Unmesh, Qixing Huang, and Karthik Ramani. Surfnet: Generating 3d shape surfaces using deep residual networks. *CoRR*, abs/1703.04079, 2017.
- [15] Shubham Tulsiani, Tinghui Zhou, Alexei A. Efros, and Jitendra Malik. Multi-view supervision for single-view reconstruction via differentiable ray consistency. In *Computer Vision and Pattern Recognition (CVPR)*, pages 209–217, 2017.
- [16] Gül Varol, Javier Romero, Xavier Martin, Naureen Mahmood, Michael J. Black, Ivan Laptev, and Cordelia Schmid. Learning from Synthetic Humans. In *CVPR*, 2017.
- [17] Daniel Vlasic, Ilya Baran, Wojciech Matusik, and Jovan Popović. Articulated mesh animation from multi-view silhouettes. In *ACM SIGGRAPH*, pages 97:1–97:9, 2008.
- [18] O. Wiles and A. Zisserman. Silnet : Single- and multi-view reconstruction by learning from silhouettes. In *British Machine Vision Conference*, 2017.
- [19] Jiajun Wu, Chengkai Zhang, Tianfan Xue, William T Freeman, and Joshua B Tenenbaum. Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling. In *Neural Information Processing Systems (NIPS)*, pages 82–90, 2016.
- [20] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *Computer Vision and Pattern Recognition (CVPR)*, pages 1912–1920, 2015.
- [21] Xinchun Yan, Jimei Yang, Ersin Yumer, Yijie Guo, and Honglak Lee. Perspective transformer nets: Learning single-view 3d object reconstruction without 3d supervision. In *Neural Information Processing Systems (NIPS)*, pages 1696–1704, 2016.
- [22] Bo Zhao, Xiao Wu, Zhi-Qi Cheng, Hao Liu, and Jiashi Feng. Multi-view image generation from a single-view. *CoRR*, abs/1704.04886, 2017. URL <http://arxiv.org/abs/1704.04886>.
- [23] Tinghui Zhou, Shubham Tulsiani, Weilun Sun, Jitendra Malik, and Alexei A. Efros. View synthesis by appearance flow. *CoRR*, abs/1605.03557, 2016. URL <http://arxiv.org/abs/1605.03557>.
- [24] Michael Zollhofer, Matthias Niessner, Shahram Izadi, Christoph Rehmann, Christopher Zach, Matthew Fisher, Chenglei Wu, Andrew Fitzgibbon, Charles Loop, Christian Theobalt, and Marc Stamminger. Real-time non-rigid reconstruction using an rgb-d camera. *ACM Transaction on Graphics*, 33(4):156:1–156:12, 2014.