

Generating Synthetic X-ray Images of a Person from the Surface Geometry

Brian Teixeira¹, Vivek Singh¹, Terrence Chen¹, Kai Ma¹, Birgi Tamersoy¹, Yifan Wu², Elena Balashova³, and Dorin Comaniciu¹

¹Medical Imaging Technologies, Siemens Medical Solutions Inc., Princeton, NJ, USA

²Temple University, Philadelphia, PA, USA

³Princeton University, Princeton, NJ, USA

Abstract

We present a novel framework that learns to predict human anatomy from body surface. Specifically, our approach generates a synthetic X-ray image of a person only from the person’s surface geometry. Furthermore, the synthetic X-ray image is parametrized and can be manipulated by adjusting a set of body markers which are also generated during the X-ray image prediction. With the proposed framework, multiple synthetic X-ray images can easily be generated by varying surface geometry. By perturbing the parameters, several additional synthetic X-ray images can be generated from the same surface geometry. As a result, our approach offers a potential to overcome the training data barrier in the medical domain. This capability is achieved by learning a pair of networks - one learns to generate the full image from the partial image and a set of parameters, and the other learns to estimate the parameters given the full image. During training, the two networks are trained iteratively such that they would converge to a solution where the predicted parameters and the full image are consistent with each other. In addition to medical data enrichment, our framework can also be used for image completion as well as anomaly detection.¹

1. Introduction

Over the past decade, there has been significant advances in realistic human body shape modeling and simulation in the graphics domain [2, 14, 25, 16], where different statistical models have been applied to learn compact parametric representations of the human body shape. However, their impact on the healthcare domain is relatively limited [21, 29]. One major reason is that existing shape

¹This feature is based on research, and is not commercially available. Due to regulatory reasons its future availability cannot be guaranteed.




Figure 1. Synthetic X-ray generated from 3D surface meshes. Each row contains data corresponding to one patient, (a) 3D surface meshes (b) corresponding ground truth images obtained using medical scanning, (c) predicted X-ray images using the proposed approach, and (d) the X-ray images with predicted markers.

modeling approaches focus primarily on the skin surface while the healthcare domain pays more attention to the internal organs. This work attempts to tackle this limitation by addressing the challenging task of estimating the internal anatomy of human body from the surface data. More specifically, our approach generates a synthetic X-ray image of a person only from the surface geometry. However the synthetic X-ray would only serve as an approximation of the true internal anatomy, we thus simultaneously predict markers which can be adjusted to update the X-ray image. Since the markers serve as spatial parameters that can be used to perturb the image, we refer to the predicted image as a *parametrized image*. Figure 1 shows examples of parametrized X-ray images generated from surface data.

Learning to predict parametrized images is a very challenging task due to strict constraints in the output space. As the training framework learns to predict images and the corresponding spatial parameters (i.e. markers), it also needs to ensure that the perturbations of these parameters lies on a manifold of ‘realistic deformations’ (e.g. realistic facial

expressions when generating face images or realistic body anatomy when generating synthetic X-ray). Since learning such output spaces (which are implicitly highly correlated) is difficult, we propose to learn a pair of networks, one trained to predict the parameters from image contents, and the other trained to predict the image contents from the parameters. When the parameters are updated, the networks are applied iteratively in a loop until convergence. To facilitate such convergent behavior during test phase, we present a novel learning algorithm that jointly learns both the networks. While some recent work have utilized predicting markers as a supplementary task within the context of multi-task learning [10, 30], to the best of our knowledge, this work is first in explicitly learning a bijection between the predicted markers and the generated images.

We use the proposed framework to generate synthetic X-ray image from 3D body surface meshes. Such a technology can be used in conjunction with existing approaches to estimate 3D body surface models from depth [21] and potentially benefit medical procedures such as patient positioning for scanning or interventional procedures. We report several impactful applications of this technology such as anomaly detection as well as completion of full X-ray from partial X-ray. We also demonstrate via experiments that the parametrized X-ray images can be used to generate training data thus helping to overcome the significant big data barrier faced during the application of the deep learning approaches in medical imaging tasks.

The key contributions of this paper can be summarized as follows:

- A novel framework to generate parametrized images using a convergent training pipeline.
- A novel technology to predict parametrized X-ray images from body surface data, which as we demonstrate can significantly impact the healthcare domain.
- Use of Wasserstein GANs (WGAN) [3] with gradient penalty method for a conditional regression task.

2. Related Work

Learning parametrized image representations can be formulated as **multi-task learning**, where one task is to generate the synthetic image from input data, and other is to ensure that the synthetic image correlates strongly with the set of control parameters. [10] presented such a multi-task architecture to regress both facial markers and a corresponding 3D face represented as a volumetric image. They present a network architecture with two stacked hourglass networks [15], one predicts the markers, and the other predicts the volumetric images taking output of the marker network as input. While they show the use of landmarks improve face prediction, there is no guarantee that the pre-

dicted markers and generated face converge to have a strong correlation, as the markers are pre-trained independently from the face network and never get any feedback from it. [18] employs an architecture where tasks are incorporated with a single network but these auxiliary tasks are categorical such as gender recognition. While such networks would have the ability to implicitly learn the correlation between the two tasks (i.e. parameter space and image space in our setting), they are difficult to train because of the large variation in losses during the training period [30].

With the advent of generative adversarial networks [5], several attempts have been made to generate synthetic training data or augment data in domain where it's either difficult to acquire or annotations are difficult to obtain. Although [22] focuses on generating data from noise, it bears similarities with our approach as it incorporates semantic parametrization during the image generation and shows that generated data can help with training deep networks. However, the sampled parameter space is limited to global scene parameters such as lighting variation, background variation etc., while in our work, we focus on the spatial parametrization which allow changing the scene structure. [32] employs GANs to learn the natural image manifold and allow user to manipulate the generated image by making structural edits, which are provided as sketches and image is updated by finding the nearest instance on the manifold. In contrast, our approach learns an explicit mapping between the image and several markers, and provides editability using these markers. [31] presents an approach for image domain transfer by learning a pair of networks with an inverse relationship, using cyclic consistency as a regularizer. However our approach learns mapping from the source domain (surface data) to target domain (X-ray image) together with control parameters (markers), and consistency is being enforced between the target image space and parameter space.

3. Parametrized Images

We refer *parametrized image* to an image that is parametrized by a set of spatially distributed markers, which can be perturbed to manipulate the contents of the image to generate realistic image variations. Such manipulation of an image via markers (spatial parameters) requires learning a bijection mapping. In this work, we achieve this by learning a pair of networks, *Marker Prediction network* (MPN), trained to predict the parameters from the image contents, and *Conditional Image Generation network* (CIGN) trained to predict the image contents given the parametrization. We use the networks to predict an initial estimate and then iteratively refine until convergence. Figure 2 shows an overview of the networks involved in the parametrized image generation. While parametrized images can be generated from noise (similar to image generation task [5]), we focus on the task of conditional image generation [9] as it naturally ap-




Figure 2. Pipeline to generate Parametrized Images. First, the synthetic X-ray image is generated from the surface data to provide an initial estimate. Next, the landmarks are detected in X-ray image, which are then in turn re-utilized to predict the X-ray image. After a few iterations, both the X-ray image and marker distribution converge, and become tightly correlated.

plies to task of generating the X-ray images from 3D body surface data.

We represent the 3D human surface mesh data with a 2-channel 2D image; the first channel stores the depth of the body surface as observed from front, and second channel stores the thickness computed by measuring the distance between the closest and furthest point as observed from front; in the rest of the document we refer to this 2 channel image as *surface image*.

3.1. Marker Prediction Network

The marker prediction network takes the surface image as well as the predicted X-ray image as input and predicts the locations for all the markers. We employ a U-Net like [19] architecture to train a network to regress from a 3-channel input image (2 surface data channels, 1 X-ray image channel) to a 17-channel heatmap image by minimizing L2-loss; these heatmaps correspond to anatomically meaningful landmarks (such as lung top, liver top, kidney center etc). Each output channel compares with the given ground truth that includes a Gaussian mask (kernel $radius = 5$, $\sigma = 1$) centered at the given target location (similar to other landmark detection approaches [28, 17, 27]).

3.2. Conditional Image Generation Network

The proposed conditional image generation network is derived from the conditional GAN architecture [9]. The generator with the U-Net architecture takes the surface image and marker heatmaps as input and outputs the synthetic X-ray image. To stabilize the CIGN training, we adopt the Wasserstein loss with gradient penalty introduced in [6], which is known to outperform other adversarial losses. The critic takes the surface image and corresponding X-ray image as input, though theoretically, a better critic model would have taken the surface images, marker maps (parameter space) and X-ray image as input to implicitly force a strong correlation between them; however fusing all the data together as a 20 channel image did not help training a useful critic in any of the GAN variants [9, 3, 6]; we noticed

that the critic was not able to utilize the marker maps to determine the real vs fake, and thus fails to provide gradients that would eventually help generating the X-ray image. The final choice of the network architecture was determined after thorough experimentation with several architectures as well as GAN variants.

4. Learning Parametrized Image Representation

In this section, we describe the procedure to train the networks (marker prediction as well as conditional image generation network). We first pre-train the networks using the available ground truth data, and subsequently refine them end-to-end to minimize the combined loss, defined as,

$$\mathcal{L} = \mathcal{L}_{MPN} + \mathcal{L}_{CIGN} \quad (1)$$

where, \mathcal{L}_{MPN} is the mean squared error between the predicted and the ground truth heat maps for the markers, \mathcal{L}_{CIGN} is loss between the predicted and ground truth X-ray image.

4.1. Pre-training Marker Prediction Network

We pre-train the marker network using the Adam optimizer [11] to minimize the MSE loss, and set the initial learning rate to 10^{-3} . During pre-training, we use the ground truth X-ray images with body surface images as input.

During the convergent training process, the input is replaced by the predicted X-ray image. This initially worsens the performance on the marker prediction network but it quickly recovers after a few epochs of convergent training, as demonstrated in the experiments.

4.2. Pre-training Conditional Image Generation Network

We pre-train the image generation network with surface images and ground truth landmark maps as input, using the Adam optimizer with initial learning rate of 10^{-3} . After

pre-training, we use the RMSProp [24] with a low learning rate of 10^{-5} . In our experiments, we found the gradient penalty variant of WGAN [6] to outperform the original WGAN with weight clipping [3]. The architecture of the critic was similar to the encoder section of the generator network. We observe that in case of WGAN, using a more complex critic helps converging faster.

During the convergent training, the network is iteratively updated using the predicted landmarks as input.

4.3. Convergent Training via Iterative Feedback

During the test phase, we apply both networks iteratively in succession until both of them reach the steady state. This implicitly requires the networks to have a high likelihood of convergence during the training stage. A stable solution sits where both the markers/parameters and synthetic image are in complete agreement with each other, suggesting a bijection. We achieve the goal by freezing one network and updating the weights of the other network using its own loss as well as the loss backpropagated from the other network. Thus, not only the networks get feedback from the ground truth, they also get feedback on how they helped each other (good markers give good X-ray image, and vice versa). The gradient flow during backpropagation is also shown in Figure 2.

The losses optimized by conditional image generation (CIGN) and marker prediction network (MPN) at each iteration are given by,

$$\mathcal{L}_{CIGN} = \mathcal{L}_{adv}(I_{gt}, I_{syn}^i) + \mathcal{L}_2(MPN(I_{syn}^i, S), M_{gt}) \quad (2)$$

$$\mathcal{L}_{MPN} = \mathcal{L}_2(M^i, M_{gt}) + \mathcal{L}_1(I_{gt}, CIGN(M^i, S)) \quad (3)$$

where, $CIGN(\cdot)$ and $MPN(\cdot)$ are deep networks depicted in functional form; I_{gt} and M_{gt} are ground truth image and markers heat maps respectively; I_{syn}^i and M^i are predicted images and markers heat maps at iteration i .

The iterative approach to train the networks to facilitate convergence is motivated by the iterative adversarial training procedure explained in [5]. However, in this case, the networks are learning to cooperate instead of compete. Similar to GAN training, there is a possibility that the training may become unstable and diverge. We address this issue by weighting the losses with appropriate scale.

While the number of epochs required to reach convergence depends on how tightly the output of the two networks correlate, in our experiments, we found 50 epochs to be sufficient. After these epochs, no significant change in X-ray or in landmarks positions were observed, suggesting convergence. Algorithm 1 details the pseudo-code for convergent training.

To validate the convergent training, we selected a random data sample from the testing set and monitored the marker displacement across iterations. Without the convergent training, the markers kept changing across iterations.

Algorithm 1 Convergent training

Require: α , the learning rate. T_s , the model generating X-ray from surface. w_m , initial MPN weights. w_t , initial CIGN weights. λ_1 and λ_2 loss scaling factors.

- 1: **while** w_m and w_t have not converged **do**
 - 2: Sample x a batch from the input data.
 - 3: Get \hat{m}, \hat{t} ground truth landmarks and X-ray
 - 4: $t_1 \leftarrow T_s(x)$
 - 5: $x' \leftarrow \text{concat}(x, t_1)$
 - 6: $m_1 \leftarrow MPN(x')$
 - 7: $e_M \leftarrow \mathcal{L}_{MPN}(m_1, \hat{m})$
 - 8: $x' \leftarrow \text{concat}(x, m_1)$
 - 9: $t_2 \leftarrow CIGN(x')$
 - 10: $e_T \leftarrow \mathcal{L}_{CIGN}(t_2, \hat{t})$
 - 11: $g_{w_m} \leftarrow \nabla_{w_m} [e_M + \lambda_1 \cdot e_T]$
 - 12: $w_m \leftarrow w_m + \alpha \cdot \text{RMSProp}(w_m, g_{w_m})$
 - 13: $x' \leftarrow \text{concat}(x, m_2)$
 - 14: $m_2 \leftarrow MPN(x')$
 - 15: $e_M \leftarrow \mathcal{L}_{CIGN}(m_2, \hat{m})$
 - 16: $g_{w_t} \leftarrow \nabla_{w_t} [e_T + \lambda_2 \cdot e_M]$
 - 17: $w_t \leftarrow w_t + \alpha \cdot \text{RMSProp}(w_t, g_{w_t})$
-

Figure 3 shows the variation of the y position of body markers over 50 iterations. Notice that with the convergent training, markers become stable after 15 iterations.

5. Experiments

To evaluate our approach, we collected 2045 full body Computed Tomography (CT) images from patients at several different hospital sites in North America and Europe. We randomly split the entire dataset into a testing set of 195 images, a validation set of 108 images and a training set with the rest. The 3D body surface meshes were obtained from the CT scans using thresholding and morphological image operations. The X-ray images were generated by orthographically projecting the CT images. All the data were normalized to a single scale using the neck and pubic symphysis body markers (since these can be easily approximated from the body surface data). All the experiments are conducted in PyTorch environment [1]. Our U-Net like networks are composed of four levels of convolution blocks (each consisting a 3 repetitions of Convolution, Batch Normalization [8] and ReLU [7]). Each network has 27 convolutional layers, with 32 filters in each layer.

5.1. Landmark Estimation

Although the purpose of the convergent training is to ensure a tight correlation between X-ray and markers, we also computed the error statistics w.r.t. the ground truth marker annotations provided by medical experts. Interestingly enough, the convergent training helped improve the




Figure 3. First row shows the evolution of all the 17 markers and second row focuses on the ‘Liver Top’ marker, which has the largest variation across people. First column presents the results before convergence and the convergent results are shown in the second column.

accuracy of the marker prediction network, though the improvement is not quantitatively significant (the mean euclidean distance dropped from 2.50 cm to 2.44 cm). An interesting observation is that the accuracy improved for some of the particularly difficult to predict markers, such as the tip of the sternum; we attribute this to the fact that convergent training facilitates more stable prediction. For example, in case of the tip of the sternum, the error dropped from 2.46 to 2.00 cm.

5.2. Synthetic X-ray Prediction

We use the widely used pix2pix [9] approach as our baseline for conditional image regression. Following [9], we use L1-loss for training generators. Using receptive field of 34×34 achieved the lowest validation error. During training, we found it difficult to optimize the network with batch size 32. We alleviated this issue by gradually reducing the batch size from 32 to 1; this enabled faster training in beginning although with blurred results but as the training continued and the batch size was reduced to 1, more details were recovered.

Table 1 shows a quantitative comparison between different approaches. Since L1 error is known to be insufficient in capturing the perceived image quality, we also report the Multi-Scale Structural Similarity index (MS-SSIM) [26] averaged over the entire testing set. Notice that the MS-SSIM score for the WGAN-GP is significantly higher than other methods. Furthermore, the convergent training is able to retain a high perceived image quality, while ensuring that the markers are strongly correlated with the X-ray image.

	L1/MAE	Multi Scale SSIM
UNet [19]	$1.62 \cdot 10^{-2}$	90.54
Pix2Pix [9]	$1.74 \cdot 10^{-2}$	90.58
WGAN [3]	$1.83 \cdot 10^{-2}$	89.04
WGAN-GP [6]	$1.651 \cdot 10^{-2}$	93.96
Conv-WGAN-GP	$1.655 \cdot 10^{-2}$	93.74

Table 1. Comparison between state-of-art algorithms on conditional image generation/X-ray prediction.

Figure 7 shows predicted X-ray images for several different surface images using these methods. Compared to ground truth, synthetic X-ray images does surprising well in certain regions such as upper thorax while does poorly in other regions such as lower abdomen where the variance is known to be significantly higher. More importantly, notice the images generated using the proposed method is sharper around organ contours as well as spine and pelvic bone structures are much clearer.

6. Applications

6.1. Training Data Generation

Due to privacy and health safety issues, medical imaging data is difficult to obtain, which creates a significant barrier for data driven analytics such as deep learning. Recently, generative deep learning model such as GANs and other variants [12, 20, 23, 4] have been effectively employed to




Figure 4. Training examples generated using parametrized images from different surface images. The annotated landmark [Left Lung Bottom] is surrounded by the dotted box.

generate realistic training data, however they offer limited control (categorical) over the nature of the perturbations.

Parametrized X-ray images offers an approach to generate medical image training data; furthermore, the spatial parametrization offers controlled perturbations such as generating data variations with lungs of certain size. For tasks such as marker detection, we further observe that since the image manifold is smooth, it’s possible to generate training data (for augmentation) together with annotations, by annotating the marker in one image and tracking it in the image domain as it is perturbed along the image manifold.

We demonstrate the effectiveness of this method on the task of detecting left lung bottom landmark (NOTE: this marker is not in the original set of 17 landmarks). We manually annotated this marker in 50 synthetic X-ray images to be used for training data. For evaluation, we manually annotated 50 ground truth X-ray images. To generate the augmented training dataset, we generate 100 random perturbations from the annotated parametrized images (by allowing the marker to move within a certain range). Since the image manifold is smooth, as the position of the marker changed in the perturbed image, we were able to propagate the annotation using coefficient normed template matching. Figure 4 shows sample synthetic training images.

We train a Fully Convolutional Network [13] to regress the marker location, depicted as a Gaussian mask, from X-ray image. We used the Adam optimizer with an initial learning rate of 10^{-3} .

To measure the usefulness of data generated using parametrized images, we created a baseline training dataset by augmenting the 50 training images using 100 random translations. Table 2 lists the error metrics as the networks are trained using the datasets for 25 epochs (after which they both overfitted). Notice that after only 5 epochs, the model trained with the parametrized training had a 0.99 cm mean error on the testing set, compared to 8.75cm for the baseline. After 25 epochs, the baseline has a mean error of 1.20 cm, while the network trained on data with parametrized perturbations has a much lower a 0.68 cm error.

Epoch	MSE (p)	MSE (t)	$\Delta(p)$	$\Delta(t)$
1	6.64.10⁻⁴	2.47.10 ⁻³	4.53	9.17
5	4.04.10⁻⁴	6.76.10 ⁻⁴	0.99	8.75
10	3.50.10⁻⁴	5.66.10 ⁻⁴	0.93	2.67
15	3.47.10⁻⁴	5.04.10 ⁻⁴	0.64	1.97
20	3.59.10⁻⁴	4.78.10 ⁻⁴	0.74	1.28
25	3.50.10⁻⁴	5.34.10 ⁻⁴	0.68	1.20

Table 2. Comparing translation training set (t) and parametrized training set (p). Metrics used are the MSE on the testing set and Δ represents the mean euclidean distance in cm.

6.2. X-ray Image Completion

As the radiation exposure is considered harmful, X-ray images are often acquired with a limited field of view, only covering a certain body region (say thorax or abdomen). Using parametric images, we can reconstruct the X-ray image of the entire body such that it is still consistent with the partial yet real image. While the reconstructed X-ray image would be of limited diagnostic use, it would be beneficial for acquisition planning in subsequent or future medical scans. Using the reconstructed X-ray, the scan region can be specified more precisely, thus potentially reducing the radiation exposure.

To reconstruct the complete X-ray, we first generate a parametrized X-ray image of the patient from the surface data. As previously mentioned, the predicted X-ray may not always correspond to the true internal anatomy. This, however, can be addressed using the markers on the parametrized image by adjusting them such that synthetic X-ray matches the real one where they overlap. Once the markers are adjusted, we regenerate the complete X-ray together with all the markers (see Figure 5). Table 3 shows the quantitative comparison between the predicted synthetic X-ray and markers, before and after being refined using the real X-ray image.

	MAE	MSE	Δ	Δ_y
Initial	1.77.10 ⁻²	1.51.10 ⁻³	2.63	2.03
Completed	1.41.10⁻²	6.60.10⁻⁴	1.46	1.07

Table 3. Using X-ray completion helped getting a more accurate X-ray and significantly reduced the landmarks error. Δ represents the mean euclidean distance and Δ_y the mean y distance. Errors are in cm.

6.3. Anomaly Detection

Another potential use case of the proposed method is anatomical anomaly detection. As the proposed method generates a representation of healthy anatomy learned from healthy patients, it can be applied for anomaly detection by

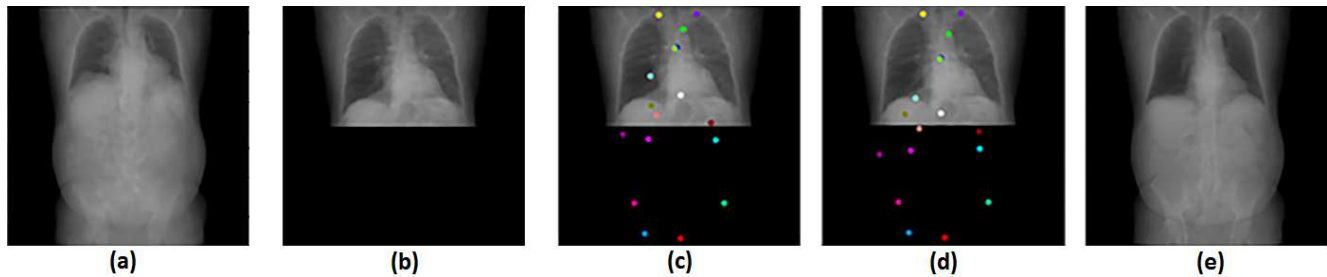


Figure 5. Synthetic X-ray completion to extrapolate an partial X-ray. (a) Predicted X-ray from surface image (b) Partial ground truth X-ray (c) Predicted markers overlaid on ground truth X-ray (d) Adjusted markers using ground truth X-ray (e) Complete synthetic X-ray that matches the ground truth X-ray where data is available.




Figure 6. Using the difference between the real X-ray (a) and the prediction (b), the anomaly can be easily detected in the resulting image (c). Here, the difference threshold was set to 10%. First row shows a missing lung detection, second row an implant detection, and last row shows the difference on a healthy patient.

quantifying the difference between the real X-ray image and the predicted one. Figure 6 illustrates such examples, where one patient has a missing lung, and the other has implant, which is sometimes overlooked by technicians. While the anatomical anomaly is easier to identify, the proposed approach with higher resolution imaging can potentially be used to suggest candidates for lung nodules (in a chest X-ray) or other pathological conditions.

7. Conclusion

We presented a novel framework that learns to predict parametrized images from partial image data, which enables natural perturbations of the predicted images. We apply

the presented method on the challenging task of predicting a synthetic X-ray image from the patient surface data, together with corresponding markers distributed over the X-ray; the predicted image can be further manipulated by adjusting the body markers while ensuring physically consistent X-ray image. The proposed technology has been demonstrated to address the significant barrier of training data scarcity in the medical domain, in addition to enabling novel use cases with benefits to the medical community such as device positioning. One of our future works is to directly predict 3D CT from the surface. This work is currently limited by the difficulties in getting sufficiently large dataset to learn the variation in 3D anatomical structure. This is also due to the fact that most of the CT scans are with limited field of view and sometimes may even have contrast depending on the scan protocol.

References

- [1] Pytorch: Tensors and dynamic neural networks in python with strong gpu acceleration. <https://github.com/pytorch/pytorch>. 4
- [2] D. Anguelov, P. Srinivasan, D. Koller, S. Thrun, J. Rodgers, and J. Davis. SCAPE: shape completion and animation of people. *ACM Trans. Graph*, 24:408–416, 2005. 1
- [3] M. Arjovsky, S. Chintala, and L. Bottou. Wasserstein generative adversarial networks. In D. Precup and Y. W. Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 214–223, International Convention Centre, Sydney, Australia, 06–11 Aug 2017. PMLR. 2, 3, 4, 5
- [4] J. Bao, D. Chen, F. Wen, H. Li, and G. Hua. CVAE-GAN: fine-grained image generation through asymmetric training. *CoRR*, abs/1703.10155, 2017. 5
- [5] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 2672–2680. Curran Associates, Inc., 2014. 2, 4




Figure 7. Comparison between X-ray images generated from surface data using various methods (best viewed on screen). The images has been contrast normalized to better appreciate the details in organ boundaries and bone structures. Notice that the images generated using the proposed method are much sharper and even the bone separation in the spine area becomes visible.

- [6] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. C. Courville. Improved training of wasserstein gans. *CoRR*, abs/1704.00028, 2017. 3, 4, 5
- [7] G. E. Hinton. Rectified linear units improve restricted boltzmann machines vinod nair, 2010. 4
- [8] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *CoRR*, abs/1502.03167, 2015. 4
- [9] P. Isola, J. Zhu, T. Zhou, and A. A. Efros. Image-to-image translation with conditional adversarial networks. *CoRR*, abs/1611.07004, 2016. 2, 3, 5
- [10] A. S. Jackson, A. Bulat, V. Argyriou, and G. Tzimiropoulos. Large pose 3d face reconstruction from a single image via direct volumetric CNN regression. *CoRR*, abs/1703.07834, 2017. 2
- [11] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. 3
- [12] D. P. Kingma and M. Welling. Auto-encoding variational bayes. *CoRR*, abs/1312.6114, 2013. 5
- [13] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. *CoRR*, abs/1411.4038, 2014. 6
- [14] M. Loper, N. Mahmood, J. Romero, G. Pons-Moll, and M. J. Black. SMPL: A skinned multi-person linear model. *ACM Trans. Graphics (Proc. SIGGRAPH Asia)*, 34(6):248:1–248:16, Oct. 2015. 1
- [15] A. Newell, K. Yang, and J. Deng. Stacked hourglass networks for human pose estimation. *CoRR*, abs/1603.06937, 2016. 2
- [16] G. Pons-Moll, J. Romero, N. Mahmood, and M. J. Black. Dyna: A model of dynamic human shape in motion. *ACM Transactions on Graphics, (Proc. SIGGRAPH)*, 34(4):120:1–120:14, Aug. 2015. 1
- [17] T. Probst, K. Maninis, A. Chhatkuli, M. Ourak, E. V. Poorten, and L. V. Gool. Automatic tool landmark detection for stereo vision in robot-assisted retinal surgery. *arXiv*, 2017. 3
- [18] R. Ranjan, V. M. Patel, and R. Chellappa. Hyperface: A deep multi-task learning framework for face detection, landmark localization, pose estimation, and gender recognition. *CoRR*, abs/1603.01249, 2016. 2
- [19] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. *CoRR*, abs/1505.04597, 2015. 3, 5
- [20] M. Rosca, B. Lakshminarayanan, D. Warde-Farley, and S. Mohamed. Variational approaches for auto-encoding generative adversarial networks. *CoRR*, abs/1706.04987, 2017. 5
- [21] V. Singh, K. Ma, B. Tamersoy, Y.-J. Chang, A. Wimmer, T. ODonnell, and T. Chen. Darwin: Deformable patient avatar representation with deep image network. In *MICCAI*, pages 497–504, 09 2017. 1, 2
- [22] L. Sixt, B. Wild, and T. Landgraf. Rendergan: Generating realistic labeled data. *CoRR*, abs/1611.01331, 2016. 2
- [23] K. Sohn, H. Lee, and X. Yan. Learning structured output representation using deep conditional generative models. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems* 28, pages 3483–3491. Curran Associates, Inc., 2015. 5
- [24] T. Tieleman and G. Hinton. Lecture 6.5—RmsProp: Divide the gradient by a running average of its recent magnitude. COURSE: Neural Networks for Machine Learning, 2012. 4
- [25] A. Tsoli, N. Mahmood, and M. J. Black. Breathing life into shape: Capturing, modeling and animating 3d human breathing. *ACM Trans. Graph.*, 33(4):52:1–52:11, July 2014. 1
- [26] Z. Wang, E. P. Simoncelli, and A. C. Bovik. Multi-scale structural similarity for image quality assessment. 2003. 5
- [27] Y. Wu and Q. Ji. Robust facial landmark detection under significant head poses and occlusion. *CoRR*, abs/1709.08127, 2017. 3
- [28] S. Yan, Z. Liu, P. Luo, S. Qiu, X. Wang, and X. Tang. Unconstrained fashion landmark detection via hierarchical recurrent transformer networks. *CoRR*, abs/1708.02044, 2017. 3
- [29] S. Y. Yeo, J. Romero, M. Loper, J. Machann, and M. Black. Shape estimation of subcutaneous adipose tissue using an articulated statistical shape model. *Computer Methods in Biomechanics and Biomedical Engineering: Imaging and Visualization*, 0(0):1–8, 2016. 1
- [30] Z. Zhang, P. Luo, C. C. Loy, and X. Tang. Facial landmark detection by deep multi-task learning. *eccv*, 2014. 2
- [31] J. Zhu, T. Park, P. Isola, and A. A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. *CoRR*, abs/1703.10593, 2017. 2
- [32] J.-Y. Zhu, P. Krähenbühl, E. Shechtman, and A. A. Efros. Generative visual manipulation on the natural image manifold. In *Proceedings of European Conference on Computer Vision (ECCV)*, 2016. 2