



**University of  
Zurich**<sup>UZH</sup>

**Zurich Open Repository and  
Archive**

University of Zurich  
University Library  
Strickhofstrasse 39  
CH-8057 Zurich  
[www.zora.uzh.ch](http://www.zora.uzh.ch)

---

Year: 2021

---

## **Micro-Ct Synthesis and Inner Ear Super Resolution via Generative Adversarial Networks and Bayesian Inference**

Li, Hongwei ; Prasad, Rameshwara G N ; Sekuboyina, Anjany ; Niu, Chen ; Bai, Siwei ; Hemmert, Werner ;  
Menze, Bjoern

DOI: <https://doi.org/10.1109/isbi48211.2021.9434061>

Posted at the Zurich Open Repository and Archive, University of Zurich

ZORA URL: <https://doi.org/10.5167/uzh-220197>

Conference or Workshop Item

Published Version

Originally published at:

Li, Hongwei; Prasad, Rameshwara G N; Sekuboyina, Anjany; Niu, Chen; Bai, Siwei; Hemmert, Werner; Menze, Bjoern (2021). Micro-Ct Synthesis and Inner Ear Super Resolution via Generative Adversarial Networks and Bayesian Inference. In: 2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI), Nice, France, 13 April 2021 - 16 April 2021. Institute of Electrical and Electronics Engineers, online.

DOI: <https://doi.org/10.1109/isbi48211.2021.9434061>

# MICRO-CT SYNTHESIS AND INNER EAR SUPER RESOLUTION VIA GENERATIVE ADVERSARIAL NETWORKS AND BAYESIAN INFERENCE

Hongwei Li<sup>\*1,2</sup>, Rameshwara G. N. Prasad<sup>\*2</sup>, Anjany Sekuboyina<sup>2</sup>, Chen Niu<sup>3</sup>, Siwei Bai<sup>4,5</sup>,  
Werner Hemmert<sup>4,5</sup>, and Bjoern Menze<sup>1,2</sup>

1. Department of Quantitative Biomedicine, University of Zurich
2. Department of Computer Science, Technical University of Munich
3. Department of Medical Imaging, First Affiliated Hospital of Xi'an Jiaotong University, China
4. Department of Electrical and Computer Engineering, Technical University of Munich
5. School of Bioengineering, Technical University of Munich

## ABSTRACT

Existing medical image super-resolution methods rely on pairs of low- and high- resolution images to learn a mapping in a fully supervised manner. However, such image pairs are often not available in clinical practice. In this paper, we address super resolution problem in a real-world scenario using unpaired data and synthesize linearly *eight times* higher resolved Micro-CT images of temporal bone structure embedded in the inner ear. We explore cycle-consistency generative adversarial networks for super-resolution and equip the model with Bayesian inference. We further introduce *Hu Moments distance* as the evaluation metric to quantify the shape of the temporal bone. We evaluate our method on a public inner ear CT dataset and have seen both visual and quantitative improvement over state-of-the-art supervised deep-learning based methods. Further, we conduct a multi-rater visual evaluation experiment and find that three inner-ear researchers consistently rate our method highest quality scores among three methods. Furthermore, we are able to quantify uncertainty in the unpaired translation task and the uncertainty map can provide structural information of the temporal bone.

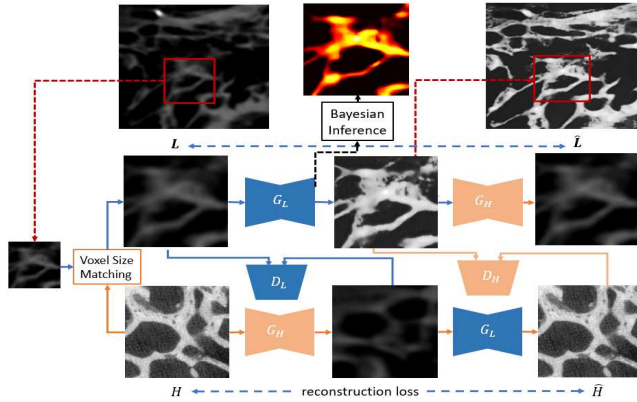
## 1. INTRODUCTION

High-quality image guided surgery has allowed clinicians to develop safe and less invasive surgical procedures despite the large inter-individual anatomical variability. Clinical computed tomography is the common imaging technique for this purpose. However, when delicate structures need to be resolved, for example the insertion of cochlear implant electrodes, a much higher resolution would be desirable. Micro computed tomography (micro-CT) can provide high resolution (HR) images that captures much more details of the human cochlear anatomy, however, it cannot be applied clinically due to the high amount of radiation dose required and due to the limited probe size which fits into the scanners. It is therefore highly clinically relevant to estimate detailed anatomical structures from low-resolution CT using super

resolution (SR) techniques.

**Related Work.** Deep learning methods have been introduced to build a mapping function from LR patches to HR ones to solve various SR tasks [12, 13, 16] and achieve superior results over traditional SR methods [2, 15, 1]. Most of the deep learning-based methods rely on pairs of LR-HR images to train in a fully supervised manner. By downsampling the HR images to generate spatially-aligned LR ones, the paired images for training the deep networks can be easily obtained. However, this strategy fails to introduce real-world characteristics from the LR domain and often struggles to generalize when given real low-resolution images. Consequently, the models trained on simulated data become less effective if there exists a large domain shift between LR and HR domains. Especially in the SR task from clinical CT (LR) to Micro-CT (HR), noise, voxel size and sensor types result in large domain gap which make it extremely difficult to generate realistic paired LR-HR images. Recently, generative adversarial networks (GAN) and their extensions [6, 20, 3] have been proposed to learn the data distribution from either paired or unpaired datasets. Although unpaired image-to-image translation using GAN was recently explored in medical imaging [10, 19], it cannot be directly applied to SR tasks because voxel sizes from the two domain are different. This work explores GAN for *real-world* CT SR tasks and quantify model uncertainty in unpaired image translation.

**Contributions.** 1) We propose a GAN-based approach to solve the unpaired super resolution problem in a clinical scenario and incorporate Bayesian inference to quantify model uncertainty to guarantee applicability in a translation to real world problems. 2) We present qualitative and quantitative results on synthesis of Micro-CT, showing the superiority of our method over existing solutions. 3) We report the results of multi-rater visual rating experiments, performed by three experienced inner-ear researchers who confirm the superior quality of generated high-resolution images while outperforming supervised deep learning approaches.



**Fig. 1:** Method overview. We upsample the input LR images by bicubic interpolation to match the voxel size. Then we sample local patches from these upsampled slices of CT scans. Then the unpaired LR and HR image patches are fed into a cycle-consistency generative adversarial network to perform the super resolution task. Finally Bayesian inference is performed to estimate the model uncertainty given a patch.

## 2. METHODS

Our super resolution approach includes two main components: 1) Unpaired image-to-image translation network and 2) Bayesian inference as shown in Figure 1.

**Problem definition.** The goal of the super resolution task is to map a low-resolution (LR) domain to a high-resolution (HR) domain. Let  $\mathcal{L}$  denote a LR image space and  $\mathcal{H}$  a HR image space. We assume that  $\mathcal{L}$  and  $\mathcal{H}$  are not paired in a common clinical setting. Considering the complexity of imaging datasets, we claim that it is difficult to introduce LR domain characteristics in medical datasets by traditional techniques such as, interpolation and noise injection [17], which restrict the applicability of SR techniques. Thus we aim to learn a direct mapping from LR domain to HR domain  $\Phi: \mathcal{L} \rightarrow \mathcal{H}$  by learning the individual underlying distributions.

**Patch-Based Unpaired Image Translation.** The super-resolution task includes two aspects: 1) Structure enhancement and 2) Image style transfer. We deploy generative adversarial networks to learn the individual underlying distributions including generator networks and discriminator networks. Since it is an unpaired image-to-image translation task, we adapt a cycle-consistency GAN inspired by [20] to handle the unpaired issue. This includes a pair of generators  $\{G_L, G_H\}$  and a pair of domain discriminators  $\{D_L, D_H\}$ . As illustrated in Figure 1, generator  $G_L$  aims to translate patches from LR to HR while generator  $G_H$  aims to translate from HR to LR. Discriminators  $\{D_L, D_H\}$  are trained to distinguish if the generated patches are real or fake in the two domains respectively. In a min-max game, the generators

try to fool the discriminators by good image translation. The detailed training losses are as follows:

*Self-reconstruction loss:*

$$\mathcal{L}_{rec} = E_{x \sim \mathcal{L}} [\|x - G_H(G_L(x))\|_1] + E_{y \sim \mathcal{H}} [\|y - G_L(G_H(y))\|_1] D_L(G_H(y)) \quad (1)$$

*Adversarial loss:*

$$\mathcal{L}_{adv} = E_{x \sim \mathcal{L}, y \sim \mathcal{H}} \{\log [D_L(x) \cdot D_L(y)]\} + E_{x \sim \mathcal{L}, y \sim \mathcal{H}} \{\log [(1 - D_H(G_L(x)))]\} + E_{x \sim \mathcal{L}, y \sim \mathcal{H}} \{\log [(1 - D_L(G_H(y)))]\} \quad (2)$$

The full objective function of our framework is:

$$\mathcal{L}_{total} = \lambda_{rec} \mathcal{L}_{rec} + \lambda_{adv} \mathcal{L}_{adv} \quad (3)$$

where  $\lambda_{rec}$  and  $\lambda_{adv}$  control the importance of self-reconstruction and adversarial training.

**Whole-slice reconstruction:** Since the method is patch-based considering the computation complexity, we further reconstruct the whole slice from individual generated image patches. As stitching patches leads to a grid effect, we eliminate this effect by first performing histogram matching for each generated patch using the corresponding LR patches and then do median filtering as a post-processing step to remove the grid effect while preserving the structural information.

**Bayesian Inference.** To estimate the model uncertainty, we employ Monte Carlo dropout [4] to run multiple forward passes through the model using a dropout layer. To derive the uncertainty for one patch  $x \in \mathcal{L}$ , we collect the predictions of  $T$  inferences with different dropout masks. Let  $G_L^{d_i}$  denote the generator  $G_L$  with dropout mask  $d_i$ . We obtain a sample of the possible model outputs distribution for a patch  $x$  as  $\{G_L^{d_0}(x), \dots, G_L^{d_T}(x)\}$ . We compute the mean and variance of  $x$ , corresponding to the mean of model posterior distribution and the estimation of model uncertainty. The predictive posterior mean  $p$  and uncertainty  $c$  are formulated as:

$$p = \frac{1}{T} \sum_{i=0}^T G_L^{d_i}; \quad c = \frac{1}{T} \sum_{i=0}^T [G_L^{d_i} - p]^2 \quad (4)$$

**Implementation.** The voxel size matching between two scanners is performed by bicubic interpolation. Each generator network contains two convolutions with a stride of two, nine residual blocks [7] and two convolutions with a stride of 1/2. Six residual blocks were used for the input with a size of  $h \times w \times 1$ , where  $h$  and  $w$  are the height and width of the input images respectively. Discriminators are convolutional neural networks for classification of real and generated images (i.e. 'real' or 'fake'). We leverage PatchGANs [8] to enforce the discriminator network to classify the local patch features to real or fake, instead of the global image. For training, we use

the Adam optimizer with a batch size of 5, a learning rate of 0.0001 and a number of 50 epochs. In all experiments, we set the hyper-parameters  $\lambda_{adv}$  and  $\lambda_{rec}$  to 1 and 10 respectively. Experiments are run on one Nvidia Titan Xp GPU. The training time is around eight hours. The inference stage takes 2 mins for 200 patches in one slice.

**Visual Evaluation Protocol and Evaluation Metrics.** Due to lack of paired images, quantitative analysis of images cannot be directly performed using traditional metrics such as PSNR and SSIM [18]. Additionally these metrics are not able to capture and evaluate useful information such as preserving the clinically relevant structures. In order to get a qualitative assessment of the images generated by our approach, we design a visual rating protocol to compare the images generated by our approach with those generated by two state-of-the-art methods (i.e. U-Net [14] and SR-GAN [9]). In each trial, the otolaryngology researchers were presented with three sets of images, with each set consisting of an original LR CT image on the left and the synthetic image generated by U-Net, SR-GAN or our approach. The rater does not know which method is used to prevent any bias. The rater is asked to rate on the following criterion: a) good resolution and low noise level and b) good shape and structural information. The otolaryngology researchers were asked to rate on a six star scale where ‘one’ represents ‘poor’ and ‘six’ represents ‘excellent’.

To quantitatively compare our method with others in the unpaired setting, we introduce *Hu Moments distance* [21] as an evaluation metric on shape similarity due to its rotation- and scale- invariance in shape encoding. However, since the generated images and HR image are not aligned, it is not possible to directly measure the shape similarity. Intuitively, the shape of the structures from real and fake images are expected to be similar. We propose *minimum Hu Moments distance* ( $m$ -HuM) to handle the similarity evaluation for unpaired slices. Given two sets of unpaired data  $\mathcal{X}$  and  $\mathcal{Y}$ , we take the minimum HuM distance for all possible combinations of the elements between the two sets:

$$m\text{-HuM} = \min\{HuM(x, y) | x \in \mathcal{X}, y \in \mathcal{Y}\} \quad (5)$$

### 3. EXPERIMENTS

**Dataset and Evaluation Metric.** We evaluate the proposed method on a public human inner ear dataset [5] including two diverse CT scanners as shown in Table 1. Specifically we perform linearly *eight times* super resolution from clinical cone beam CT to the state-of-the-art Micro CT. For pre-processing, all the low-resolution scans are matched in voxel size using bicubic interpolation and for each slice, we crop the background to obtain a region of interest by thresholding and using a bounding box. Then they are normalized using  $z$ -score normalization and the intensity is then normalized to the range of [0, 1] to meet the requirement of the network output. We then extract patches with a size of  $256 \times 256$  from

whole slices with a stride of 128 pixels resulting in a training set of 5K patches. We use six LR scans from subset ‘AL\_1’ and six HR ones from subset ‘BH\_1’ as a training set. Also, we use one paired scan (with same patient ID) from ‘BL\_1’ and ‘AH\_1’ as a test set. Note that even though the two scans are from the same patient, it is extremely difficult to align them spatially.

**Qualitative and Quantitative Results.** We firstly evaluate our method on the test set with state-of-the-art methods U-Net [14] and SR-GAN [11]. Note that both of them are supervised-learning methods and require paired data for training. For these two methods, we synthesize LR images by bicubic interpolation considering the voxel size difference and inject Gaussian noise as used in previous SR tasks [17]. As shown in Fig. 2, visually our method enhances the temporal bone structure and preserves good-quality shape information compared to other supervised approaches. We also observe that SR-GAN generates some artefacts in both background and structure boundary by hallucinating noisy patches which are dominated by continuous partial volume effects that are of no relevance for the bony structures, while U-NET preserves the shape structure but struggles to enhance the local structure. Although SR-GAN uses a GAN component to enhance the image quality, it does not guarantee the model can be generalized to real data. We further perform  $T$  forward passes through the model using a dropout layer to collect the predictions. In practice, we set  $N = 100$  considering the computation complexity. The heatmaps in Fig. 2 depict the uncertainty obtained for the input patches. We observe that uncertainty maps match well with the temporal bone structure and thus they can serve as segmentation maps to interpret the structure in LR images.


Table 2 shows the quantitative comparison of three methods using  $m$ -HuM. For each whole slice, we calculate the  $m$ -HuM by comparing it to the whole test set. The results of all testing slices are averaged. It demonstrates that our method outperforms U-Net and SR-GAN in terms of preserving the shape of temporal bone structure.

**Visual Evaluations.** In order to assess how our approach generates images with clinical relevance, three raters with a mean of 7+ years of experience rated the image quality based on the criterion mentioned in Section 2. Results on the test set for the different metrics (i.e. resolution, noise level and preservation of shape and structure) are presented in the boxplots in Fig. 3. It shows that our method outperforms U-Net and SR-GAN in both metrics. We observe that U-Net produces higher quality images than SR-GAN in terms of resolution and noise level. For each sample, we conduct Wilcoxon rank-sum test on the paired rating scores on two testing scans from three raters to compare the methods. Results show that the pair of metric ratings on our approach with other two methods are significantly different (with p-values  $< 0.0001$ ). This demonstrates that



Scanner	Voxel Size (mm <sup>3</sup> )	Volume Size	Subjects
Micro-CT	0.018 × 0.018 × 0.018	1828 × 1828 × 1828	7
cone beam CT	0.15 × 0.15 × 0.15	668 × 668 × 668	7

**Table 1:** Data characteristics of a subset of the public inner ear dataset including one Micro-CT and one clinical CT. We used five cone beam CT scans and five Micro-CT scans for the unpaired training and the rest for testing.



**Fig. 2:** Qualitative results of three approaches. The region marked with red rectangle in the first row shows the artefacts generated by SR-GAN. The region marked with yellow circle shows that our approach generates good shape and local structure with less noise. The uncertainty maps highlight the structure of the temporal bone and can serve as a rough segmentation map.

Methods	U-Net [14]	SR-GAN [9]	Ours
$m\text{-HuM} \downarrow$	3.12	1.29	<b>0.28</b>

**Table 2:** Comparison with state-of-the-art supervised methods using *minimum Hu Moments distance*.  $\downarrow$  indicates that a lower value represents a higher shape similarity.


our method significantly outperforms U-Net and SR-GAN in image quality. As claimed in Section 2, our method benefits from learning the real distribution of LR domain instead of introducing artificial factors. This strategy strengthens its applicability to real data.

#### 4. SUMMARY AND CONCLUSION

This work introduces a super resolution approach integrating generative adversarial network and Bayesian inference. For the first time, we address Micro-CT image synthesis in a real-world scenario using only unpaired data. This approach is potentially helpful for image-guided surgery of cochlear interventions. As a promising approach, we will investigate its usage for safer and less invasive surgical procedures.

#### 5. COMPLIANCE WITH ETHICAL STANDARDS


The study was carried out in accordance with *the Declaration of Helsinki*.



**Fig. 3:** Rating results by three inner-ear researchers. Our method generates good-quality images and significantly outperforms other supervised methods, benefiting from learning distributions on real data using GAN.

#### 6. ACKNOWLEDGEMENT

The project was funded in part by the German Research Foundation (DFG HE6713/2-1).



**Fig. 4:** One synthetic high-resolution micro-CT slice (right) generated from a clinical CT image (left) of one testing subject. Notably it is very challenging to match the corresponding ground truth high-resolution image due to the random 3D orientation.

## 7. REFERENCES

- [1] D. C. Alexander, D. Zikic, A. Ghosh, R. Tanno, V. Wotschel, J. Zhang, E. Kaden, T. B. Dyrby, S. N. Sotiropoulos, H. Zhang, et al. Image quality transfer and applications in diffusion mri. *NeuroImage*, 152:283–298, 2017.
- [2] K. Bahrami, F. Shi, I. Rekić, Y. Gao, and D. Shen. 7t-guided super-resolution of 3t mri. *Medical physics*, 44(5):1661–1677, 2017.
- [3] Y. Choi and et al. Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. In *CVPR*, pages 8789–8797, 2018.
- [4] Y. Gal and Z. Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. *arXiv preprint arXiv:1506.02142*, 2015.
- [5] N. Gerber, M. Reyes, L. Barazzetti, H. M. Kjer, S. Vera, M. Stauber, P. Mistrik, M. Ceresa, N. Mangado, W. Wimmer, et al. A multiscale imaging and modelling dataset of the human inner ear. *Scientific data*, 4:170132, 2017.
- [6] I. Goodfellow. Nips 2016 tutorial: Generative adversarial networks. *arXiv preprint arXiv:1701.00160*, 2016.
- [7] K. He and et al. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016.
- [8] J. Johnson and et al. Perceptual losses for real-time style transfer and super-resolution. In *ECCV*, pages 694–711. Springer, 2016.
- [9] C. Ledig, L. Theis, F. Huszár, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4681–4690, 2017.
- [10] H. Li, J. C. Paetzold, A. Sekuboyina, F. Kofler, J. Zhang, J. S. Kirschke, B. Wiestler, and B. Menze. Diamondgan: Unified multi-modal generative adversarial networks for mri sequences synthesis. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 795–803. Springer, 2019.
- [11] M.-Y. Liu, T. Breuel, and J. Kautz. Unsupervised image-to-image translation networks. In *Advances in neural information processing systems*, pages 700–708, 2017.
- [12] O. Oktay, W. Bai, M. Lee, R. Guerrero, K. Kamnitsas, J. Caballero, A. de Marvao, S. Cook, D. O’Regan, and D. Rueckert. Multi-input cardiac image super-resolution using convolutional neural networks. In *International conference on medical image computing and computer-assisted intervention*, pages 246–254. Springer, 2016.
- [13] J. Park, D. Hwang, K. Y. Kim, S. K. Kang, Y. K. Kim, and J. S. Lee. Computed tomography super-resolution using deep convolutional neural network. *Physics in Medicine & Biology*, 63(14):145011, 2018.
- [14] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [15] W. Shi, J. Caballero, C. Ledig, X. Zhuang, W. Bai, K. Bhatia, A. M. S. M. de Marvao, T. Dawes, D. O’Regan, and D. Rueckert. Cardiac image super-resolution with global correspondence using multi-atlas patchmatch. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 9–16. Springer, 2013.
- [16] R. Tanno, D. E. Worrall, A. Ghosh, E. Kaden, S. N. Sotiropoulos, A. Criminisi, and D. C. Alexander. Bayesian image quality transfer with cnns: exploring uncertainty in dmri super-resolution. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 611–619. Springer, 2017.
- [17] Z. Wang, J. Chen, and S. C. Hoi. Deep learning for image super-resolution: A survey. *arXiv preprint arXiv:1902.06068*, 2019.
- [18] P. Welander, S. Karlsson, and A. Eklund. Generative adversarial networks for image-to-image translation on multi-contrast mr images—a comparison of cyclegan and unit. *arXiv preprint arXiv:1806.07777*, 2018.
- [19] M. Yurt, S. U. H. Dar, A. Erdem, E. Erdem, and T. Çukur. mustgan: Multi-stream generative adversarial networks for mr image synthesis. *arXiv preprint arXiv:1909.11504*, 2019.
- [20] J.-Y. Zhu and et al. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *CVPR*, pages 2223–2232, 2017.
- [21] J. Žunić, K. Hirota, and P. L. Rosin. A hu moment invariant as a shape circularity measure. *Pattern Recognition*, 43(1):47–57, 2010.