

# NEURAL RADIANCE PROJECTION

Pham Ngoc Huy<sup>†</sup>, and Tran Minh Quan<sup>†‡</sup>  
 Email : {hpham, qtran}@talosix.com

<sup>†</sup>Talosix, Ho Chi Minh City, Vietnam

<sup>‡</sup>VinUniversity, Ha Noi, Vietnam

## ABSTRACT

The proposed method, Neural Radiance Projection (NeRP), addresses the three most fundamental shortages of training such a convolutional neural network on X-ray image segmentation: dealing with missing/limited human-annotated datasets; ambiguity on the per-pixel label; and the imbalance across positive- and negative- classes distribution. By harnessing a generative adversarial network, we can synthesize a massive amount of physics-based X-ray images, so-called Variationally Reconstructed Radiographs (VRRs), alongside their segmentation from more accurate labeled 3D Computed Tomography data. As a result, VRRs present more faithfully than other projection methods in terms of photo-realistic metrics. Adding outputs from NeRP also surpasses the vanilla UNet models trained on the same pairs of X-ray images.

*Index Terms*— GAN, Chest X-ray, NeRF, NeRP

## 1. INTRODUCTION

X-ray (XR) imaging can create photos of the body inside. The images show the parts in different intensities of grayscale. It is because various tissues absorb particular amounts of radiation. The most typical use of X-rays is checking for broken bones (fracture), but X-rays are also used for other purposes: chest XRs can spot pneumonia, or mammograms use XRs to look for breast cancer. The amount of radiation from an XR is small. For example, a chest XR gives out a radiation dose similar to the amount our body naturally exposes to from the environment over ten days.

Although XR imaging is widely used as an indispensable tool in diagnostic medicine thanks to its fast and cheap operations, the collected XR datasets are often long-tail and imbalanced, i.e., the number of negative cases is dominant compared to those that have disease-detection, because many typical cases take regular healthcare checkup annually. There is also a worldwide shortage of fully understanding the image interpretation since many overlapping things are projected on the same pixels. In order to resolve this issue, more advanced imaging techniques such as Computed Tomography (CT) or Magnetic Resonance (MR) are solicited to disentangle the ambiguity in 3D space. They provide more reliable and golden ground truth to examine the patients. However,

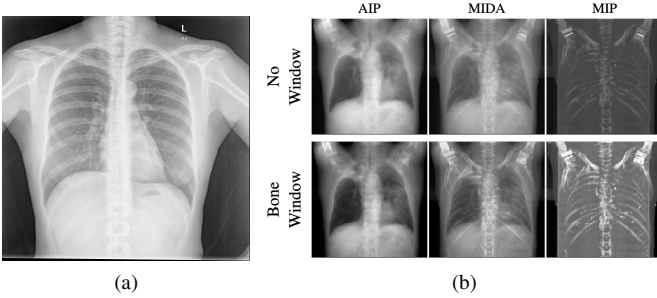
only severely positive cases need to take further scans due to their expense in cost and time operations. On the other hand, Artificial Intelligence and Machine Learning have shown remarkable performance recently in the automated evaluation of medical images. Nevertheless, they require well-annotated, balanced, and relatively large datasets to produce an excellent supervised-learning model for deployment. These drawbacks pose significant challenges in starting an XR-based medical image analysis study for new pathology problems when the requirements from both sides burden the others.

Our work stems from a recently advanced solution on Neural Radiance Field (NeRF) [1] that attempts to synthesize/reconstruct the scene from sparse-view images. We leverage the ray marching and ray sampling strategy to implement our differentiable Neural Radiance Projection (NeRP) to deal with common issues which hinder the deep learning approach on direct and end-to-end solutions for radiograph image analysis: **(1)** Missing/Limited amount of data: Usually, when we start a brand new image analysis problem with a specific goal, we are entirely lacking the data or having only a few samples for particular tasks. **(2)** Annotations often come with uncertainties: Due to the blurry effect and depth estimation, it is ambiguous to label the per-pixel region of interest, especially on projected radiographs when many things can accumulate on the same pixels. **(3)** Imbalanced data since positively severe cases rarely appear in the datasets: in our case, less positive images for certain classes in XR collection can be readjusted by more positive CT scans.

In summary, the proposed method aims to produce many more physics-based XR images with golden CT scan labels. Our contributions are several-fold as follows:

- Variational Cameras within proximities to generate a large number of labeled pairs of XR images.
- A Generative-Adversarial Network (GAN) to enhance the realism of generated XR data.
- A data augmentation to enrich the limited/missing training pair of XR images.

To the best of our knowledge, this work is the first attempt that leverages heterogeneous data from both CT and XR to train a segmentation model and improve its performance compared to the baselines of using XR data only. We demonstrate its usability in real-world applications of segmenting X-ray images.



**Fig. 1:** (a) A real XR image (b) XR-like images generated by DDR methods with or without pre-windowing.

## 2. METHOD

### 2.1. XR-like image generation

The most direct way to transform a typical 3D CT volume data to a 2D image is to average the intensity (ray accumulation) along particular directions (view direction) to generate such an X-ray-like image. This approach gives us a brief sense of how overall the 3D CT data looks like in the context of projection, but it does not look similar to the real X-ray image that is commonly taken, in terms of contrast or intensity, even with common pre-windowing setups of CT data (see Fig. 1). It can be explained that this averaging intensity projection (AIP) does not fully reflect the photon attenuation of heterogeneous absorber (CT data) measuring in the Hounsfield Unit (HU) scale. Several techniques such as MIDA [2] combines a linear lookup table (intensity versus opacity) and non-linear ones such as Maximum-Intensity Projection (MIP) to promote better certain material visualization.

The 2D Digitally Reconstructed Radiographs (DRRs) generated from 3D CT images by more advanced methods such as volume rendering algorithms [3, 4], have better representations of photo-realistic projections. There is also seminal work that attempts to generate physics-based XR such as DeepDRR [5], XraySyn [6]. They both require hand-crafted designed matters of interest (masks) to determine material thickness criteria to discriminate different radiant interactions. On the other hand, XPGAN [7] produces non-physics-based XR with source location-free. However, even if it can approximate the projection by a 3D-2D neural network, the novel view synthesis is not amendable since the viewpoint does not involve in the parameters of training. In an extreme case, slightly rotating the CT volume results in a very different XR image. Therefore, it motivates us to develop a fully end-to-end learning method to create a photo that looks realistic to the particular XR images without human assistance of segmentation or prior knowledge of radiance interaction but still be able to manipulate the viewpoint.

### 2.2. Model overview

As shown in Fig. 2, the proposed method, NeRP, consists of these building blocks: a Generator  $\mathcal{G}$  includes a Mapper  $\mathcal{M}$ , a Viewer  $\mathcal{V}$ ; and a Discriminator  $\mathcal{D}$  that are parameterized by following objective functions:

$$\min_{\theta_{\mathcal{G}}} \max_{\theta_{\mathcal{D}}} L_{adv}(\mathcal{G}, \mathcal{D}) + \lambda L_{reg}(\mathcal{G}) \quad (1)$$

$$\min_{\theta_{\mathcal{M}}, \theta_{\mathcal{V}}} \max_{\theta_{\mathcal{D}}} L_{adv}(\mathcal{M}, \mathcal{V}, \mathcal{D}) + \lambda L_{reg}(\mathcal{M}, \mathcal{V}) \quad (2)$$

The Mapper  $\mathcal{M}$  takes a single channel 3D CT volume  $y_{1 \times D \times H \times W} \sim P_Y$ , which has been normalized to produce a two-channel volume  $y_{2 \times D \times H \times W}^{\mu\gamma}$  that consists of a Matter Field ( $y_{1 \times D \times H \times W}^{\mu}$  volume) combined with an Opacity Field ( $y_{1 \times D \times H \times W}^{\gamma}$  volume). The subscripts indicate the shapes of tensors which have channel  $C = \{1, 2\}$ , depth  $D$ , height  $H$ , and width  $W$ , respectively. They will be omitted for simplicity. The maximum bandwidth of  $y^{\gamma}$  is set **100-fold smaller** than  $y^{\mu}$  to foster the XR-like effect and discriminate different roles between them. Otherwise, these two channels are interchangeable and behave similarly to each other. Then the Viewer  $\mathcal{V}$  accepts the radiance  $y^{\mu\gamma}$ , constructs the beam emission absorption model with ray bundle  $r \sim P_R$  so that these rays can interact with the implicit volume function  $y^{\mu}$  and the opacity  $y^{\gamma}$ . The proposed method also randomly samples a (variationally) perspective camera pose  $c \sim P_C$  within a certain proximity of Field of View (FoV) to form the final image on the screen. The result is therefore called Variationally Reconstructed Radiographs (VRRs), as distinguish to DRRs. We employ Generative Adversarial Networks (GANs) [8] in XPGAN [7] to treat the VRR results as fake images and train the networks adversarially with the true XR samples that drawn from the real datasets  $x \sim P_X$ . The adversarial loss  $L_{adv}$  is formally defined in Eq. 3 with the expectation of NeRP result is expanded in Eq. 4:

$$L_{adv}(\mathcal{G}, \mathcal{D}) = \mathbb{E}_{x \sim P_X} [\log \mathcal{D}(x)] + \mathbb{E}_{y \sim P_Y} [1 - \log \mathcal{D}(\mathcal{G}(y))] \quad (3)$$

$$\mathbb{E}_{y \sim P_Y} [1 - \log \mathcal{D}(\mathcal{G}(y))] = \mathbb{E}_{\substack{y \sim P_Y \\ r \sim P_R \\ c \sim P_C}} [1 - \log \mathcal{D}(\mathcal{V}(\mathcal{M}(y), r, c))] \quad (4)$$

Last but not least, we add an  $\ell_1$  per-pixel loss between the generated image and the ray tracing rendered version, weighted by  $\lambda$  to regularize the projected images to have the same anatomical structure as of its original CT volumes.

$$L_{reg}(\mathcal{G}) = \mathbb{E}_{y \sim P_Y} [\|\mathcal{G}(y) - Proj(y)\|_1] \quad (5)$$

For wide-range choices of  $\{0.05, 0.1, 0.5, 1, 5, 10\}$ ,  $\lambda = 1$  yields better convergence and image quality compared to other options.

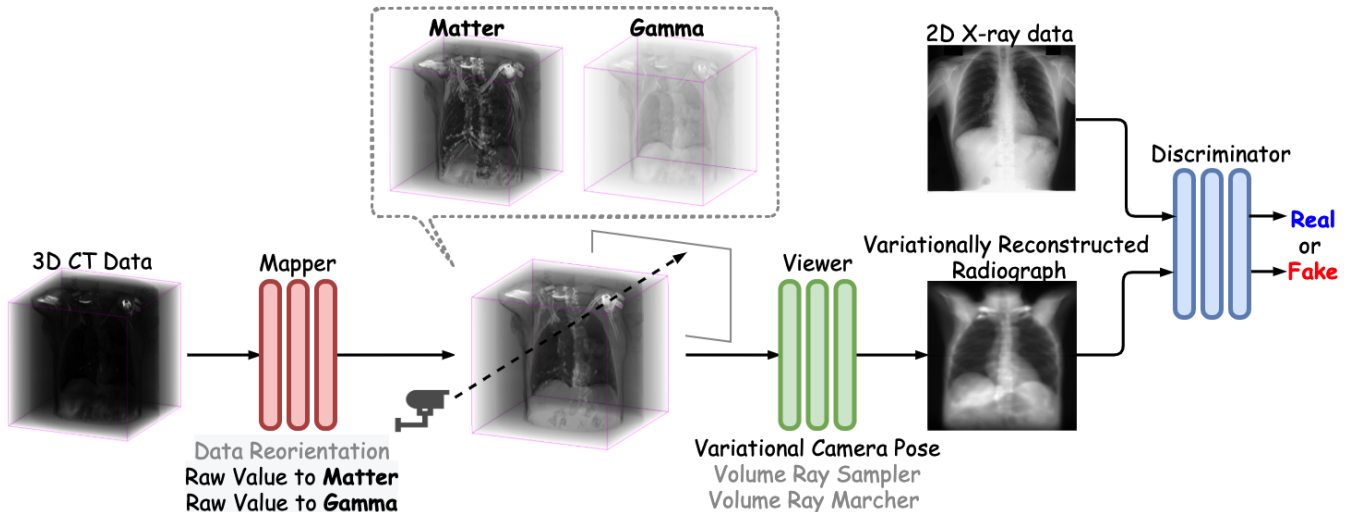


Fig. 2: Neural Radiance Projection Overview.

### 2.3. Implementation details

To leverage the power of state-of-art models, we utilize architectures of the current latest methods. Specifically, we use V-Net [9] for Mapper  $\mathcal{M}$  and U-Net [10] for Discriminator  $\mathcal{D}$  [11]. It is worth noting that having U-Net as a Discriminator stabilizes the training process. For the Viewer  $\mathcal{V}$ , the ray bundle is set as the same as image size ( $256 \times 256$ ), each ray is sampled with 512 points. The entire system is trained for 200 epochs with an NVIDIA 3080 Ti GPU for one week.

## 3. DATA

Table 1: Statistics of CT volumes and XR images.

Type	Dataset	Description	Total
CT	NSCLC [12]	Lung Cancer	402
	MosMed [13]	COVID-19	1110
	Total		<b>1512</b>
CXR	ChinaSet [14]	Tuberculosis	566
	Montgomery [14]	Tuberculosis	138
	JSRT [15]	Lung nodules	247
	Total		<b>951</b>

The CT volumes were prepared from two datasets. The first is Non-Small Cell Lung Cancer CT Volume, comes from the radio-genomic database of non-small cell lung cancer patients at Stanford University Medical Center and Palo Alto Veterans Affairs Healthcare System. This collection has 402 annotated CT from a cohort of 211 patients. The MosMed dataset is comprised of COVID-19 related CT scans from the hospitals in Moscow, Russia. Besides that, the chest X-Ray images come from three different sources. Both Montgomery and ChinaSet were publicized by the US National Library of Medicine in 2014. The former dataset comes from the Tuberculosis screening program of Montgomery County, Mary-

land, USA, and consists of 138 anterior-posterior CXR images; the latter contains 566 frontal projections in both normal and tuberculosis patients [14]. The Japanese Society of Radiological Technology published the JSRT database in 1998. This source includes 247 chest radiographs from Japanese and American institutions. These images were used to identify the pulmonary nodules in terms of five degrees [15]. The dataset statistics are collected in the Table 1.

## 4. RESULTS

### 4.1. Quality of generated images

Table 2: Image quality metrics:

Higher IS and lower FID, KID mean better image quality.

Method	IS $\uparrow$	FID $\downarrow$	KID $\downarrow$
DRR [3]	$1.799 \pm 0.050$	197.071	$0.226 \pm 0.013$
DeepDRR [5]	$2.323 \pm 0.062$	<b>153.399</b>	<b><math>0.158 \pm 0.013</math></b>
NeRP	<b><math>2.454 \pm 0.072</math></b>	161.675	$0.161 \pm 0.013$

Many relevant metrics can be used to evaluate the results on how photo-realistic they are over the real datasets: such as Inception Score (IS) [16], Frechet Inception Distance (FID) score [17] and Kernel Inception Distance (KID) score [18]. We pre-generate 4020 XR images from 402 CT data from NLCSC (one volume produces 10 projections with different viewpoints) for each setup (DRR, DeepDRR, NeRP) and compare against each other over the distribution of Chest XR datasets [14, 15]. Table 2 shows the results of each metrics. Our method has similar scores to DeepDRR, and both are better than DRR. Note that although DeepDRR has more parameters to control, such as spectrum, photon counts, or scattering, it does not support generating the projected labels.

## 4.2. Lung Segmentation

### 4.2.1. Missing annotated XR data

In order to show the advantages of NeRP, we compare its performance against DRRs on Lung segmentation of Chest XR in the scenario of lacking the training data entirely, only test on given manual sets as discussed. To make a fair assessment, we set the same viewpoint-generated VRRs from NeRP and DRRs. The only difference between them is that our radiance  $y^{\mu\gamma}$  is learnable while DRRs have fixed HU-based  $y^\mu$  and constant  $y^\gamma$ . Fig. 3 indicates some VRR samples generated by NeRP that have been used for training a vanilla UNet to perform pixel-based lung segmentation problems. As can be seen in Table 3, using our method can outperform a DRR-based dataset up to 5% on the Dice Score [19] metrics. This can be explained that the XR images produced by NeRP are much typical than by DRRs, and hence require less overhead transfer compared to those direct DRR CT-generated samples.

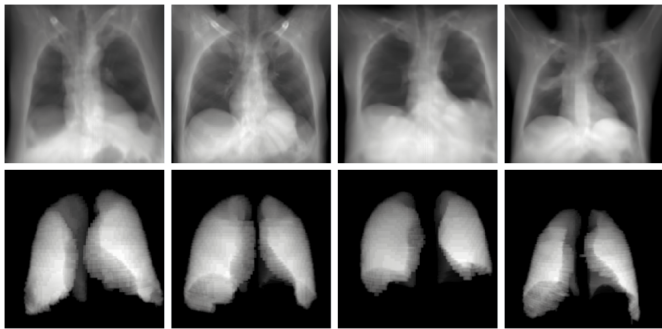


Fig. 3: Samples of NeRP pairs for Lung Segmentation task.

Table 3: Comparison of Dice Scores of DRR- and NeRP-generated pairs from 3D CT data, and test on human-annotated 2D XR images for Lung Segmentation.

Dataset	ChinaSet	Montgomery	JSRT	Average
DRR [3]	0.8033	0.8060	0.8229	0.8068
NeRP	<b>0.8491</b>	<b>0.8261</b>	<b>0.8357</b>	<b>0.8412</b>

### 4.2.2. Limited annotated XR data

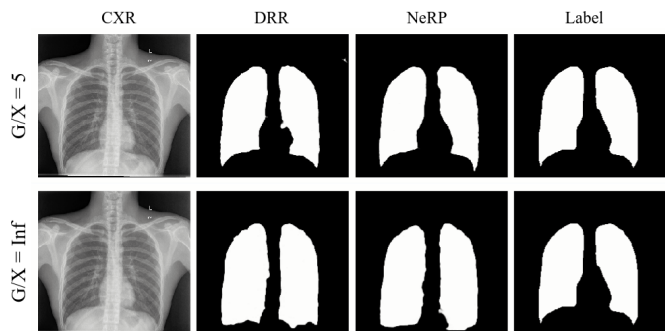


Fig. 4: Samples of NeRP data for Lung Segmentation task.

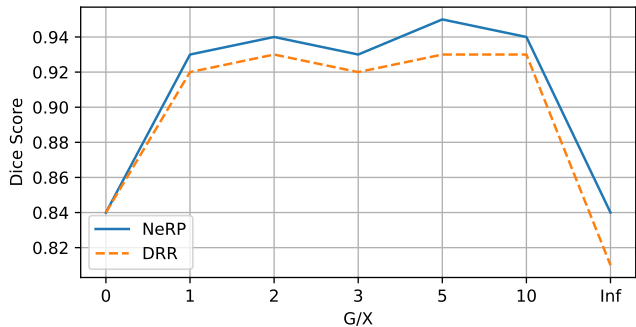


Fig. 5: Dice scores of Lung Segmentation task.

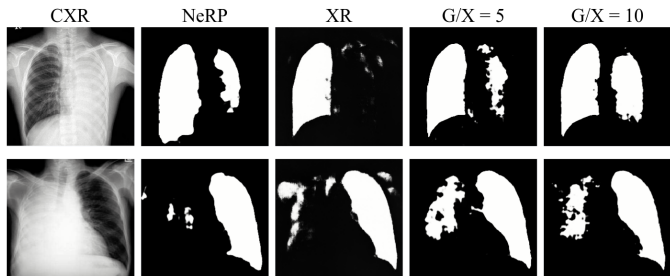


Fig. 6: Comparison in Pleural effusion lung collapse samples.

In this assessment, we gradually include portions of Chest XR with Lung annotation to train the U-Net and observe the performance over the quota of having extra annotated subsets of 2D XR data (see Fig. 4). Fig. 5 shows the Dice Scores on the separated test set for this setup. The horizontal axis represents the ratio ( $G/X$ ) between the number of the generated images and XR images: 0 means that the model was trained from XR only, and  $Inf$  means that model was trained from the generated images only. The results indicate that mixing the handful of annotated data with NeRP images yields better performance (up to 2%) over conventional DRR.

Furthermore, we choose extreme cases of Chest XR images with Pleural Effusion Lung Collapse, i.e., lung lobe on one side is filled by high contrast and brightness pixel values. As depicted in Fig. 6, combining NeRP and XR data promotes better lung regions even if the appearance of dark areas is partially missing.

## 5. CONCLUSION

We present NeRP, which can produce VRRs. Comparing to conventional DRRs, NeRP can generate much more faithful images in common photo-realistic metrics. This work addresses the three most fundamental issues of establishing a machine learning model to solve XR-based image analysis problem: (1) Missing/Limited amount of data; (2) Uncertain annotation; and (3) Imbalanced datasets. By leveraging the ability of massive amount pairs of segmentation data created by NeRP, we can utilize more golden ground truth and ready-to-use CT data to balance the currently long-tail datasets with more accurate projected 3D labels. This usefulness will be applied for more XR-based image analysis problems such as rare cancer detection, diagnosis, and prognosis.

## 6. COMPLIANCE WITH ETHICAL STANDARDS

All procedures performed in studies involving human participants were in accordance with the ethical standards of the institutional and/or national research committee and with the 1964 Helsinki declaration and its later amendments or comparable ethical standards. For this type of retrospective study, formal consent is not required.

## 7. CONFLICTS OF INTEREST

The authors have no conflicts of interest to disclose.

## 8. REFERENCES

- [1] B. Mildenhall, P.P. Srinivasan, M. Tancik, J.T. Barron, R. Ramamoorthi, and R. Ng, “NeRF: Representing scenes as neural radiance fields for view synthesis,” in *Proc. of ECCV*, 2020.
- [2] Z. Hussain, F. Gimenez, D. Yi, and D. Rubin, “Differential data augmentation techniques for medical imaging classification tasks,” in *AMIA annual symposium proceedings*. American Medical Informatics Association, 2017, vol. 2017, p. 979.
- [3] Robert L. Siddon, “Fast calculation of the exact radiological path for a three-dimensional CT array,” *Medical Physics*, vol. 12, no. 2, pp. 252–255, Mar. 1985.
- [4] F. Jacobs, E. Sundermann, B. De Sutter, M. Christiaens, and I. Lemahieu, “A fast algorithm to calculate the exact radiological path through a pixel or voxel space,” *CIT. Journal of computing and information technology*, vol. 6, no. 1, pp. 89–94, 1998.
- [5] M. Unberath, J.N. Zaech, S. C. Lee, B. Bier, J. Fotouhi, M. Armand, and N. Navab, “DeepDRR—A Catalyst for Machine Learning in Fluoroscopy-guided Procedures,” in *Proc. of MICCAI*, 2018.
- [6] Cheng Peng, Haofu Liao, Gina Wong, Jiebo Luo, S. Kevin Zhou, and Rama Chellappa, “Xraysyn: Realistic view synthesis from a single radiograph through ct prior,” in *Proc. of the 35th AAAI Conference*, 2021.
- [7] T.M. Quan, H.M. Thanh, T.D. Huy, N.D.T Chanh, N.T.P. Anh, P.H. Vu, N.H. Nam, T.Q. Tuong, V.M. Dien, B.V. Giang, B. Trung, and S.Q.H. Truong, “XPGAN: XR projected GAN for improving COVID-19 image classification,” in *Proc. of IEEE ISBI*, 2021, pp. 1509–1513.
- [8] Ian Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, “Generative adversarial nets,” *Advances in neural information processing systems*, vol. 27, 2014.
- [9] F. Milletari, N. Navab, and S.-A. Ahmadi, “V-net: Fully convolutional neural networks for volumetric medical image segmentation,” in *Proc. of 3DV*. IEEE, 2016, pp. 565–571.
- [10] Olaf Ronneberger, Philipp Fischer, and Thomas Brox, “U-net: Convolutional networks for biomedical image segmentation,” in *Proc. of MICCAI*, 2015, pp. 234–241.
- [11] Edgar Schonfeld, Bernt Schiele, and Anna Khoreva, “A u-net based discriminator for generative adversarial networks,” in *Proc. of IEEE CVPR*, 2020, pp. 8207–8216.
- [12] S. Bakr, O. Gevaert, S. Echeagaray, K. Ayers, M. Zhou, M. Shafiq, H. Zheng, J. A. Benson, W. Zhang, A. N. C. Leung, M. Kadoch, C. D. Hoang, J. Shrager, A. Quon, D. L. Rubin, S. K. Plevritis, and S. Napel, “A radio-genomic dataset of non-small cell lung cancer,” *Sci Data*, vol. 5, pp. 180202, 10 2018.
- [13] S.P. Morozov, A.E. Andreychenko, N.A. Pavlov, A.V. Vladzimirskyy, N.V. Ledikhova, V.A. Gombolevskiy, I.A. Blokhin, P.B. Gelezhe, A.V. Gonchar, and V. Yu Chernina, “Chest CT scans with COVID-19 related findings dataset,” *preprint arXiv:2005.06465*, 2020.
- [14] S. Jaeger, S. Candemir, S. Antani, Y. X. Wáng, P. X. Lu, and G. Thoma, “Two public CXR datasets for computer-aided screening of pulmonary diseases,” *Quant Imaging Med Surg*, vol. 4, no. 6, pp. 475–477, Dec 2014.
- [15] J. Shiraishi, S. Katsuragawa, J. Ikezoe, T. Matsumoto, T. Kobayashi, K. Komatsu, M. Matsui, H. Fujita, Y. Kodera, and K. Doi, “Development of a digital image database for chest radiographs with and without a lung nodule: receiver operating characteristic analysis of radiologists’ detection of pulmonary nodules,” *AJR Am J Roentgenol*, vol. 174, no. 1, pp. 71–74, Jan 2000.
- [16] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen, “Improved techniques for training GANs,” *Advances in neural info. proc. systems*, vol. 29, pp. 2234–2242, 2016.
- [17] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter, “GANs trained by a two time-scale update rule converge to a local nash equilibrium,” *Adv. in neural info. proc. sys.*, vol. 30, 2017.
- [18] Mikołaj Bińkowski, Danica J Sutherland, Michael Arbel, and Arthur Gretton, “Demystifying MMD gans,” *preprint arXiv:1801.01401*, 2018.
- [19] K. H. Zou, S. K. Warfield, A. Bharatha, C. M. Tempany, M. R. Kaus, S. J. Haker, W. M. Wells, F. A. Jolesz, and R. Kikinis, “Statistical validation of image segmentation quality based on a spatial overlap index,” *Acad Radiol*, vol. 11, no. 2, pp. 178–189, Feb 2004.