

UCLA

UCLA Previously Published Works

Title

USING A GENERATIVE ADVERSARIAL NETWORK FOR CT NORMALIZATION AND ITS IMPACT ON RADIOMIC FEATURES

Permalink

<https://escholarship.org/uc/item/5ms7n1gd>

Authors

Wei, Leihao
Lin, Yannan
Hsu, William

Publication Date

2020

Peer reviewed

© 2020 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

USING A GENERATIVE ADVERSARIAL NETWORK FOR CT NORMALIZATION AND ITS IMPACT ON RADIOMIC FEATURES

Leihao Wei^{1,3} Yannan Lin^{2,3} William Hsu^{2,3,4}

¹ Electrical and Computer Engineering, University of California, Los Angeles ² Bioengineering, University of California, Los Angeles

³ Medical & Imaging Informatics, University of California, Los Angeles ⁴ Radiological Sciences, University of California, Los Angeles

ABSTRACT

Computer-Aided-Diagnosis (CADx) systems assist radiologists with identifying and classifying potentially malignant pulmonary nodules on chest CT scans using morphology and texture-based (radiomic) features. However, radiomic features are sensitive to differences in acquisitions due to variations in dose levels and slice thickness. This study investigates the feasibility of generating a normalized scan from heterogeneous CT scans as input. We obtained projection data from 40 low-dose chest CT scans, simulating acquisitions at 10%, 25% and 50% dose and reconstructing the scans at 1.0mm and 2.0mm slice thickness. A 3D generative adversarial network (GAN) was used to simultaneously normalize reduced dose, thick slice (2.0mm) images to normal dose (100%), thinner slice (1.0mm) images. We evaluated the normalized image quality using peak signal-to-noise ratio (PSNR), structural similarity index (SSIM) and Learned Perceptual Image Patch Similarity (LPIPS). Our GAN improved perceptual similarity by 35%, compared to a baseline CNN method. Our analysis also shows that the GAN-based approach led to a significantly smaller error (p-value < 0.05) in nine studied radiomic features. These results indicated that GANs could be used to normalize heterogeneous CT images and reduce the variability in radiomic feature values.

Index Terms— lung cancer, radiomics, generative adversarial networks, deep neural networks, denoising

1. INTRODUCTION

Lung cancer is the leading cause of cancer death in the United States. The National Lung Screening Trial (NLST) clinical trial demonstrated a 20% mortality rate reduction in patients who underwent chest low-dose computed tomography (LDCT) [1]. Radiologists routinely interpret images acquired from different scanners at varying dose levels and slice thicknesses. These differences in acquisition affect morphology and texture-based features that are used to describe pulmonary nodules, leading to inconsistencies in the detection and characterization of lesions in images. Our goal is to develop a method for normalizing heterogeneous CT scans to generate scans with common reconstruction

parameters on which computer-aided diagnosis (CADx) systems can be executed. Several prior works have used deep learning to denoise LDCT images. Chen et al. [2] used a residual encoder-decoder convolutional neural network to optimize mean square error (MSE) loss, while Wolterink[3] and Yang[4] used a generative adversarial network (GAN) to reconstruct normal dose CT images with better textures. Differences in dose and slice thickness affect downstream image analysis such as radiomic feature generation and may cause CADx systems to give inconsistent results. Prior studies have shown that automated detection and nodule segmentation performance is impacted due to heterogeneous CT images[5, 6]. Therefore, to facilitate cross-platform/protocol CT image analysis, reducing the error/variability in volume-wise CT appearance due to differences in reconstruction parameters and protocols is of special interest. One approach to addressing differences in resolution, for example, is to apply super-resolution techniques, such as asking a network to reconstruct the high-resolution image. You et al.[7] showed that a CycleGAN[8] was capable of recovering down-sampled images.

We introduce a novel CT image normalization method that utilizes a 3D GAN with spectral-norm. The contributions of our work are as follows:

1. Using a 3D GAN model, we normalize the dose and slice thickness of a CT scan simultaneously.
2. In addition to traditional metrics such as PSNR and SSIM, we quantitatively assess perceptual quality using Learned Perceptual Image Patch Similarity (LPIPS)[9].
3. We show that radiomic feature variability when comparing the original heterogeneous acquisition and our normalized reconstructions have been significantly reduced.

2. METHOD

In prior work, GANs have been used to reconstruct photorealistic images for enhancement, denoising, and super-resolution tasks [10, 11]. Here, we developed and evaluated a 3D GAN as the basis for volumetric normalization of CT scans. Our model has a unique spectral-norm layer to further improve




Fig. 1. Network structure of our GAN-based approach

training stability and speed. For comparison, we also implemented a baseline model based on a convolutional neural network (CNN) with MSE loss. Each mapping between the original and normalized dose and slice thickness was trained as an independent model. We calculated image quality using peak signal-to-noise ratio (PSNR), structural similarity (SSIM) and Learned Perceptual Image Patch Similarity (LPIPS) along the axial, coronal, and sagittal views. We extracted radiomic features of a nodule in the test set and computed the absolute error between the original reduced dose/thick slice image, the CNN- and GAN-based outputs, and features generated from the normal dose/thin slice image.

2.1. Dataset

Our dataset consists of raw projection data extracted from Siemens CT scanners (Sensation 64, Definition Flash) for 40 patients who underwent chest LDCT exams. Most datasets introduce Poisson noise to the post-reconstructed image to simulate reduced doses. However, we utilize a previously validated physics-based model [12] that injects noise into the sinogram of each scan to simulate reduced dose levels at 10%, 25% and 50% of the original LDCT acquisition. The final CT images were reconstructed via filtered back projection using a medium kernel at both 1.0mm and 2.0mm slice thicknesses with an image size of 512×512 . We partitioned data into 30 training cases, 5 validation cases, and 5 testing cases.

2.2. Network Architecture

GANs consist of a generator G and a discriminator D . The generator maps an input volume x to a volume representing a standardized acquisition $G(x)$ with the reference being y . The discriminator D is trained using both the reconstructed standardized scans and real reference scans as inputs to differentiate between $G(x)$ and y . D constantly judges the similarity between $G(x)$ and y to improve the performance of the generator. A good generative model is achieved when the dis-

criminator can no longer distinguish between the $G(x)$ and y . However, GANs are notoriously difficult to train. To improve training stability, Wasserstein GAN-GP [13] imposes local regularization on the discriminator to satisfy the Lipschitz continuity constraint by penalizing the gradients. Here, we use spectral-norm [14] to satisfy the same constraint. Spectral-norm is a robust global regularization technique that divides every weight matrix by its largest singular value, as opposed to calculating a 2nd order gradient penalty, which can be computationally expensive. Inspired by Enhanced Deep Residual Networks (EDSR) [15], our generator contains multiple layers of residual blocks to extract features followed by an up-sampling block in the longitudinal dimension. This up-sampling block effectively increases z resolution, mitigating the partial-volume effect in thick slice scans. We use hinge loss for the discriminator to restrict D to focus on hard samples that are difficult to differentiate. The generator loss function contains a $\mathcal{L}_1$ content loss and an adversarial loss. Loss functions $V_D(G, D)$ and $V_G(G, D)$ are shown in equations 1 and 2, and training proceeds with alternating D and G updates, $\min_G \max_D [V_D(G, D) + V_G(G, D)]$, where Θ and W are network parameters. The network structures are summarized in figure 1. For comparison, we trained the same generator network using only the $\mathcal{L}_1$ content loss as our baseline CNN model.

$$V_D(G, D) = \mathbb{E}_{y \sim p_y} [\min(0, -1 + D_\Theta(y))] + \mathbb{E}_{x \sim q_x} [\min(0, -1 - D_\Theta(G_W(x)))] \quad (1)$$

$$V_G(G, D) = -\alpha_1 \mathbb{E}_{x \sim q_x} [D_\Theta(G_W(x))] + \alpha_2 \mathbb{E}_{\substack{x \sim q_x \\ y \sim p_y}} \|G_W(x) - y\|_1 \quad (2)$$

2.3. Model Training and Tuning

In the training stage, we inputted randomly generated patches of size $16 \times 64 \times 64$ (depth, height, width), excluding patches that were primarily outside of the body. Each voxel value was scaled from Hounsfield units to $[0, 1]$. The outputted patch from the network has a dimension of $32 \times 64 \times 64$. We trained three different GANs to accommodate the following normalization scenarios: 2.0mm slice thickness at 10% dose, 25% dose, and 50% dose to 1.0mm slice thickness at 100% dose. The learning rate of both the generator and discriminator is set to $1e - 5$. Based on hyperparameters used in [14], D/G update ratio was set to 1. We used Adam optimizer with $\beta_1 = 0.5$ and $\beta_2 = 0.999$. For the generator loss function weights, we used $\alpha_1 = 1, \alpha_2 = 5e - 3$. These values were set using a grid search on the validation set. The batch size was set to 14, and we trained the model using 60k iterations on an NVIDIA Titan XP GPU, taking 60 hours. During inference, we used half-float precision on an input volume of $512 \times 512 \times 32$ to save GPU memory. The final inference outputs are stitched back together with an overlap of 4 voxels in z




Fig. 2. Normalization results. Each row represents the axial, coronal and sagittal view of an ROI that contains a nodule. Target images are shown in column a followed by the input images (columns b, e, h), the CNN-based (columns c, f, i) and the GAN-based (columns d, g, j) results generated from the input.

direction. Processing time for the whole volume of a patient’s scan with 160 slices is around 45 seconds.

Table 1. Image quality assessment. Axial (Ax), coronal (Co), and sagittal (Sa) views results are shown in each metric’s sub-rows. LPIPS measures the perceptual similarity between the target and CNN-based, GAN-based normalized images (the **lower** the better). Dose-slice pairs normalization scenarios: **A.** 10% dose, 2.0mm slice to 100% dose, 1.0mm slice **B.** 25% dose, 2.0mm slice to 100% dose, 1.0mm slice **C.** 50% dose, 2.0mm slice to 100% dose, 1.0mm slice

Metric		A		B		C	
		CNN	GAN	CNN	GAN	CNN	GAN
PSNR(dB)	Ax	29.94	28.30	30.12	28.85	30.69	29.63
	Co	32.04	30.34	32.17	30.90	32.78	31.65
	Sa	30.19	28.55	30.34	29.08	30.91	29.84
SSIM	Ax	0.7294	0.6780	0.7380	0.7027	0.7646	0.7372
	Co	0.7196	0.6655	0.7288	0.6904	0.7552	0.7249
	Sa	0.7270	0.6762	0.7357	0.7009	0.7619	0.7342
LPIPS	Ax	0.3090	0.2147	0.2649	0.1992	0.2202	0.1844
	Co	0.3286	0.2081	0.2862	0.1970	0.2514	0.1829
	Sa	0.3059	0.1937	0.2672	0.1817	0.2342	0.1642

2.4. Image Quality Assessment

Many metrics have been proposed to assess image reconstruction quality. Among them, PSNR and SSIM have been widely used to measure local differences to a reference image. However, these functions are computed based on low-level features. Optimizing loss corresponding to these functions (e.g., using mean-squared error) leads to overly smoothed images with reduced texture. To better assess the image quality, we used LPIPS, a perceptual metric that utilizes a pre-trained

VGG network [9] to generate similarity scores of high-level semantic features between two images. A lower LPIPS value represents a closer distance to the reference/target image. For each metric, results were calculated along the axial (x-y), coronal (x-z), and sagittal (y-z) planes. Results were averaged over all 5 scans in the test set. The same calculations were performed on all three normalization scenarios.

2.5. Radiomic Feature Error Analysis

Due to heterogeneous CT acquisitions, radiomic features extracted from images can vary widely. To evaluate the impact of our GAN-based normalization approach on radiomic feature values, we used `pyradiomics` [16] to extract nine representative radiomic features. To mimic a CAD system making a diagnosis on a nodule whether being benign or malignant, we only focused on a region of interest (ROI) which occupied a patch with a dimension of $30 \times 64 \times 64$. For each slice i in all three normalization scenarios, we calculated a feature value $\hat{x}_i$ after normalization and the reference feature value x_i . The normalized error on i th slice was defined as $\frac{|\hat{x}_i - x_i|}{x_i}$. The same analysis was repeated for all nine features. To find the significant differences in extracted features, statistical tests were also performed as paired Wilcoxon signed-rank sum test with a significance threshold at $p < 0.05$.

3. RESULTS

Normalization results are shown in figure 2. The ROI displayed contains a part-solid lung nodule (solid component




Fig. 3. Radiomic feature errors for raw, CNN-, and GAN-based outputs across all three normalization scenarios. Radiomic features where the GAN achieved a significantly lower error compared to the CNN are denoted with an asterisk. (Wilcoxon signed rank test, $p < 0.05$). A representative ROI from which radiomic features were calculated is shown in the top right corner.

< 5mm) with poorly circumscribed margins. The displayed lung window is centered at -600 HU with a width of 1500 HU. The GAN achieved better perceptual quality than CNN with richer textures and sharp edges. Enhancement is more pronounced in the coronal and sagittal views.

3.1. Image Quality

Table 1 summarizes image quality results. Our GAN method achieves superior LPIPS. On average, compared to the baseline CNN, our GAN improved perceptual similarity by 35%, 29%, and 25% for each of the three normalization scenarios. The largest improvement was achieved when the input was 2.0mm slice and 10% dose (the extreme normalization scenario). However, PSNR and SSIM were lower than that of the CNN method. Although the CNN achieves higher PSNR and SSIM, the resulting images appear to lose important textures that are useful for a radiologist or CADx algorithm to inform a diagnosis.

3.2. Radiomics Features

While the proposed GAN-based approach yielded superior image perceptual quality, it might also generate artifacts that were not real (e.g., unrealistic textures within the nodule). Figure 3 shows box plots that summarize the effect of the GAN model on the error of radiomic features (e.g., comparing the value of the normalized scan to that of the standard reference scan). Among all 9 radiomic features, the GAN-based method has significantly lower mean feature errors compared to the non-normalized, where the largest p-value is 0.035. The GAN-based method also has a lower error compared to CNN-generated scans in 8 features marked with an asterisk.

4. DISCUSSION

By normalizing heterogeneous CT scans to a standardized dose and slice thickness, we show that the GAN-based method not only leads to a better perceptual appearance in the normalized images but also results in reduced variability in radiomics features. These results are potentially valuable for developing CADx systems that can achieve more consistent detection and classification performance when presented with images acquired using different CT hardware platforms/protocols. In addition, different from prior work which only focused on denoising or super-resolution independently, our method performs these tasks simultaneously.

Some limitations of this work include: 1. We have not evaluated how normalization would impact nodule detection and classification tasks by CADx. Since the ultimate goal is to improve the CADx system performance by normalizing scans from different protocols or reconstruction parameters, we need to further evaluate the impact of the GAN-based approach to these specific diagnosis tasks. We intend to assess the impact of our normalization approach on the performance of existing systems, including our own [17]. 2. Since reconstruction kernel also plays an important role in image textures and thus impacts CADx performance on texture-based feature extraction, we also plan to investigate the feasibility of normalizing different kernels. 3. We have not yet explored how to tune the GAN to faithfully generate features within and immediately surrounding the nodule, which are the critical areas to enhance for the target classification task.

5. ACKNOWLEDGEMENTS

The authors acknowledge funding support for this work from the National Cancer Institute of the National Institutes of Health under award R01 CA210360 and the National Science Foundation under grant #1722516. We thank Drs. Michael McNitt-Gray and John Hoffman for providing access to their physics-based simulated chest CT dataset. Computing resources were provided by a donation of a Titan Xp GPU by NVIDIA Corporation. The content is solely the responsibility of the authors and does not necessarily represent the official views of sponsor agencies.

6. REFERENCES

- [1] National Lung Screening Trial Research Team, “Reduced lung-cancer mortality with low-dose computed tomographic screening,” *New England Journal of Medicine*, vol. 365, no. 5, pp. 395–409, 2011.
- [2] Hu Chen, Yi Zhang, Mannudeep K Kalra, Feng Lin, Yang Chen, Peixi Liao, Jiliu Zhou, and Ge Wang, “Low-dose ct with a residual encoder-decoder convolutional neural network,” *IEEE transactions on medical imaging*, vol. 36, no. 12, pp. 2524–2535, 2017.
- [3] Jelmer M Wolterink, Tim Leiner, Max A Viergever, and Ivana Išgum, “Generative adversarial networks for noise reduction in low-dose ct,” *IEEE transactions on medical imaging*, vol. 36, no. 12, pp. 2536–2545, 2017.
- [4] Q. Yang, P. Yan, Y. Zhang, H. Yu, Y. Shi, X. Mou, M. K. Kalra, Y. Zhang, L. Sun, and G. Wang, “Low-dose ct image denoising using a generative adversarial network with wasserstein distance and perceptual loss,” *IEEE Transactions on Medical Imaging*, vol. 37, no. 6, pp. 1348–1357, June 2018.
- [5] M Wasi Wahi-Anwar, Nastaran Emaminejad, John Hoffman, Grace H Kim, Matthew S Brown, and Michael F McNitt-Gray, “Towards quantitative imaging: stability of fully automated nodule segmentation across varied dose levels and reconstruction parameters in a low-dose ct screening patient cohort,” in *Medical Imaging 2018: Computer-Aided Diagnosis*. International Society for Optics and Photonics, 2018, vol. 10575, p. 105751L.
- [6] Nastaran Emaminejad, Pechin Lo, Shahnaz Ghahremani, Grace H Kim, Matthew S Brown, and Michael F McNitt-Gray, “The effects of slice thickness and radiation dose level variations on computer-aided diagnosis (cad) nodule detection performance in pediatric chest ct scans,” in *Medical Imaging 2017: Computer-Aided Diagnosis*. International Society for Optics and Photonics, 2017, vol. 10134, p. 101340B.
- [7] Chenyu You, Guang Li, Yi Zhang, Xiaoliu Zhang, Hongming Shan, Mengzhou Li, Shenghong Ju, Zhen Zhao, Zhuiyang Zhang, Wenxiang Cong, et al., “Ct super-resolution gan constrained by the identical, residual, and cycle learning ensemble (gan-circle),” *IEEE Transactions on Medical Imaging*, 2019.
- [8] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros, “Unpaired image-to-image translation using cycle-consistent adversarial networks,” in *Computer Vision (ICCV), 2017 IEEE International Conference on*, 2017.
- [9] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang, “The unreasonable effectiveness of deep features as a perceptual metric,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 586–595.
- [10] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio, “Generative adversarial nets,” in *Advances in neural information processing systems*, 2014, pp. 2672–2680.
- [11] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan 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*, 2017, pp. 4681–4690.
- [12] S Young, P Lo, J Hoffman, M Wahi-Anwar, F Noo, M Brown, and M McNitt-Gray, “Th-ab-207a-05: A fully-automated pipeline for generating ct images across a range of doses and reconstruction methods,” *Medical physics*, vol. 43, no. 6Part44, pp. 3860–3860, 2016.
- [13] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville, “Improved training of wasserstein gans,” in *Advances in neural information processing systems*, 2017, pp. 5767–5777.
- [14] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida, “Spectral normalization for generative adversarial networks,” *arXiv preprint arXiv:1802.05957*, 2018.
- [15] Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee, “Enhanced deep residual networks for single image super-resolution,” in *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 2017, pp. 136–144.
- [16] Joost JM Van Griethuysen, Andriy Fedorov, Chintan Parmar, Ahmed Hosny, Nicole Aucoin, Vivek Narayan, Regina GH Beets-Tan, Jean-Christophe Fillion-Robin, Steve Pieper, and Hugo JWL Aerts, “Computational radiomics system to decode the radiographic phenotype,” *Cancer research*, vol. 77, no. 21, pp. e104–e107, 2017.
- [17] Yannan Lin, Leihao Wei, Simon X. Han, Denise R. Aberle, and William Hsu, “EDICNet: an end-to-end detection and interpretable classification network for pulmonary nodules on computed tomography,” in *Medical Imaging 2020: Computer-Aided Diagnosis*. 2020, SPIE.