

Unsupervised Detection of Lung Nodules in Chest Radiography Using Generative Adversarial Networks

Nitish Bhatt^{1,*}, David Ramón Prados¹, Nedim Hodzic¹, Christos Karanassios¹, and H.R. Tizhoosh²

Abstract—Lung nodules are commonly missed in chest radiographs. We propose and evaluate P-AnoGAN, an unsupervised anomaly detection approach for lung nodules in radiographs. P-AnoGAN modifies the fast anomaly detection generative adversarial network (f-AnoGAN) by utilizing a progressive GAN and a convolutional encoder-decoder-encoder pipeline. Model training uses only unlabelled healthy lung patches extracted from the Indiana University Chest X-Ray Collection. External validation and testing are performed using healthy and unhealthy patches extracted from the ChestX-ray14 and Japanese Society for Radiological Technology datasets, respectively. Our model robustly identifies patches containing lung nodules in external validation and test data with ROC-AUC of 91.17% and 87.89%, respectively. These results show unsupervised methods may be useful in challenging tasks such as lung nodule detection in radiographs.

Clinical relevance— P-AnoGAN can detect challenging lung nodules without need for labelled training data. Such designs could be appealing for software solutions for diagnosis and triage in high volume or high acuity settings.

I. INTRODUCTION

Lung cancer represents the leading cause of cancer deaths worldwide claiming an estimated 1.76 million lives in 2018 [1]. Chest radiography (CXR) remains widely used for initial radiological examination in patients with suspected lung cancer. However, the National Lung Screening Trial showed that CXR has low sensitivity (73.5%) and high specificity (91.3%) in lung cancer detection [2]. Many factors contribute to this low sensitivity including observer errors in scanning CXR images, recognition of nodules, and decisions regarding the clinical relevance of subtle opacities [3]. Highly variable nodule characteristics such as size and location add complexity to this task. Computer aided diagnosis (CAD) methods have been shown to improve lung nodule detection performance of radiologists in CXR [4].

Recently, deep learning CAD tools using Convolutional Neural Networks (CNN) for lung cancer detection have been proposed and validated [5]. However, these methods have all applied supervised learning. Supervised deep learning methods face several challenges in medical imaging including the high costs of obtaining large volumes of labelled training data, inaccuracies in data labelling due to intra-observer variability, lack of representativeness due to labelling by one

or just a few experts, and difficulty learning the nodule class due to clinical imbalance in healthy and unhealthy data. Unsupervised anomaly detection methods can circumvent many of these challenges by learning instances of a healthy (normal) class and then distinguishing instances of unhealthy data deviating from the learned normal class.

Several previous studies have proposed unsupervised anomaly detection models based on generative adversarial network (GAN) models. Schlegl et al. introduced the well-known f-AnoGAN model by first training a Wasserstein GAN (WGAN) to learn the distribution of normal data and then using a deep convolutional encoder to learn the optimal mapping between images and the GAN latent space to detect abnormalities in optical coherence tomography (OCT) [6]. Furthermore, Sun et al. reported an abnormal-to-normal translation GAN for unsupervised detection and segmentation of lesions in brain MRI [7]. Very recently, Nakao et al. presented an auto-encoding GAN model to classify lung opacities in CXR [8]. However, quantitative results of model performance for lung nodules specifically were not reported. To our understanding, the use of deep anomaly detection techniques for lung nodules in chest radiographs has not systematically been studied. Moreover, previous methods have not validated or tested unsupervised anomaly detection methods using “external datasets” and thus cannot shed light on the robustness and generalizability of such methods in medical imaging.

In this paper, we investigate the feasibility of unsupervised anomaly detection for lung nodules in chest radiography. Our contributions can be summarized as follows: (1) performing automated detection of lung nodules in chest radiographs without the pathology labels; (2) extending the popular f-AnoGAN method by using a progressive growing GAN, proposing new robust loss for encoder training, and comparing ResNet and DenseNet encoders; and (3) analyzing our network using external validation and testing datasets.

II. METHODS

A. Datasets

We utilized three datasets for training, validation, and testing obtained from independent sources. The training dataset for our method contained only healthy data instances. The Indiana University CXR collection consists of 8,121 images obtained from 3,996 patients [9]. De-identified radiology reports with findings by a board certified radiologist were also available for each patient. We manually reviewed radiological reports and included 320 consecutive patients with no acute cardiopulmonary findings in our training dataset.

Research supported by Faculty of Engineering, University of Waterloo
¹N. Bhatt, D. Ramón Prados, N. Hodzic, C. Karanassios are with the Department of Systems Design Engineering, University of Waterloo, Waterloo, ON, N2L 3G1, Canada

²H. R. Tizhoosh is the director of KIMIA Lab, University of Waterloo, Waterloo, ON N2L 3G1, Canada; tizhoosh@uwaterloo.ca

*Corresponding author email: n4bhatt@uwaterloo.ca

Our validation dataset consisted of patients with lung nodules (i.e., *abnormality*) and healthy patients obtained from the National Institutes of Health (NIH) ChestX-ray14 dataset [10]. The NIH dataset contains 108,948 images obtained from 32,717 patients, however the majority of these images were automatically annotated using a natural language processing labeller which is error prone and thus could not be used for evaluation. The dataset contained a small subset of unhealthy patients with specific pathologies annotated using a bounding box by a board certified radiologist. From this subset, we included all 79 patients with lung nodules in our validation dataset. Since no manually labelled healthy patients were available, we randomly sampled 25 automatically annotated patients with “no findings” label. A board certified physician was consulted, and 10 patients confirmed to have no cardiopulmonary findings or medical support devices were included in our validation dataset.

The testing dataset consisted of healthy patients and patients with lung nodules from the Japanese Society of Radiological Technology (JSRT) database [11]. The JSRT database contains 93 healthy patients and 154 patients with lung nodules. All images were manually labelled by using consensus of a panel of board-certified radiologists and nodule location was provided using image coordinates. Therefore, images in this dataset were suitable for model testing. In all datasets, only frontal (posteroanterior) views were used, and all images were resized to 1024×1024 pixels.

B. Data Preprocessing

Contrast Limited Adaptive Histogram Equalization (CLAHE) was used to enhance visual local image contrast (kernel size= 128; clip limit= 0.01) [12]. An open-source, robust pre-trained deep network based on criss-cross attention (CCNet) was used to perform image segmentation and extract lung regions of interest (ROI) [13]. Patients with nodules outside of the lung ROI were excluded (6 validation patients and 13 test patients). As a result, our validation and test datasets contained 73 and 141 patients with nodules, respectively. Uniform patches of size 64×64 were extracted from the lung ROIs. All patches extracted from healthy patients were labelled as normal. All patches from unhealthy patients which intersected the nodule centroid were labelled as abnormal. Thus, we obtained normal training patches ($n = 70798$); normal ($n = 9010$) and abnormal ($n = 1062$) validation patches; and normal ($n = 96765$) and abnormal ($n = 2105$) test patches. Model training and evaluation were performed using these extracted patches.

C. Model Architecture

Our model (P-AnoGAN) consisted of a GAN and encoder network (Fig. 1) inspired by the f-AnoGAN method [6]. We utilized the Progressive GAN (PGAN) which consisted of two mirrored CNNs: generator G and discriminator D . The generator maps a random 512-point latent vector to a 64×64 image patch using 3×3 convolution layers, upsampling, leaky ReLU activation, and per-pixel normalization of feature vectors. Simultaneously, the discriminator learns




Fig. 1. Model with Encoder (E), Generator (G), and Discriminator (D).

classification of real and fake images using 3×3 convolution layer, downsampling, and leaky ReLU [14]. Our encoder was based on the ResNet-50 and DenseNet-169 architectures which embedded input image patch \mathbf{x} to a 512-point latent vector \mathbf{z} . The trained GAN generator and encoder were combined in an encoder-decoder-encoder architecture where the encoder performed the mapping $\mathbf{z} = E(\mathbf{x})$ and generator acted as a decoder to reconstruct $\tilde{\mathbf{x}} = G(\mathbf{z})$. The encoder then mapped $\tilde{\mathbf{x}}$ to the latent space $\tilde{\mathbf{z}} = E(\tilde{\mathbf{x}})$.

D. Model Training

P-AnoGAN was implemented using Python (version 3.7) and PyTorch (version 1.8.1) and training was completed on a single Tesla V100 GPU (NVIDIA Corporation, Santa Clara, CA) with CUDA version 11.2. For model training, we used unlabeled patches $\mathbf{x} \in X$, where X is the manifold of patches containing healthy anatomy from our training dataset.

The PGAN was trained to model the variability in X . The generator learned the mapping between a latent space $Z \in \mathbb{R}^{512}$ to X by generating “fake” images $\tilde{\mathbf{x}}$ from randomly sampled latent vector \mathbf{z} . Adam optimizer with $\beta_1 = 0.0$, $\beta_2 = 0.99$, base learning rate of 0.001 was used as in [14]. Mini-batch size of 16 was set. Training was completed for 48,000 iterations for the 4×4 scale; 96,000 iterations for 8×8 , 16×16 , and 32×32 ; and 200,000 iterations for the 64×64 scale. The generator and discriminator networks were trained progressively at growing resolution scale by upsampling and adding network blocks. Training started at the 4×4 scale. Initially, the generator consisted of a single network block comprised of convolutional layers and ReLU activation which outputted 4×4 “fake” patches and the discriminator also comprised of a single block comparing real and generated patches. After iterations were completed, we moved to the next 8×8 scale by upsampling and adding network blocks. Linear fading was performed whereby output features from 4×4 blocks were scaled by a parameter α and concatenated to the output of the 8×8 blocks; the α decremented linearly by $1/600$ every 32 iterations. This process was repeated until the final scale.

After the completion of GAN training, generator and discriminator networks weights were fixed. Next, the ResNet-50 and DenseNet-169 encoders were trained. Encoders learned the mapping from manifold X to latent space Z such that provided a healthy image \mathbf{x} , the encoded latent vector could be used by the generator to accurately reconstruct the input image $\tilde{\mathbf{x}} = G(E(\mathbf{x}))$. Several loss functions were tested for encoder training. Three loss functions izi , ziz , and izi_f were

introduced in Schlegl et al. [6]. The image-latent-image (*izi*) loss was calculated as the pixel-level mean square error (MSE) between original and reconstructed patches as

$$L_{izi} = \frac{1}{n} \|\mathbf{x} - \tilde{\mathbf{x}}\|^2, \quad (1)$$

where n is the number of pixels in the image patch. To use latent-image-latent (*ziz*) loss function in training, a random latent vector $\mathbf{z} \in Z$ was sampled and used to generate an image $\tilde{\mathbf{x}} = G(\mathbf{z})$ which was subsequently re-encoded to a latent $\tilde{\mathbf{z}} = E(\tilde{\mathbf{x}})$. The *ziz* loss was then the MSE between original and re-encoded latent vectors as

$$L_{ziz} = \frac{1}{d} \|\mathbf{z} - \tilde{\mathbf{z}}\|^2, \quad (2)$$

where d is the number of points in the latent vector. The image-latent-image loss with discriminator guidance (*izif*) was similar to *izi* but included a loss term quantifying MSE between features from an intermediate layer of the discriminator for the original \mathbf{D}_f and reconstructed $\tilde{\mathbf{D}}_f$ images, as

$$L_{izif} = \frac{1}{n} \|\mathbf{x} - \tilde{\mathbf{x}}\|^2 + \frac{\gamma}{n_d} \|\mathbf{D}_f - \tilde{\mathbf{D}}_f\|^2, \quad (3)$$

where n_d is the number of discriminator features and γ is a weighting factor. Finally, we proposed the combination of *ziz* and *izif* losses to enforce similarity between images and latent vectors. The encoder network was used to embed the reconstructed image to a latent vector $\tilde{\mathbf{z}} = E(\tilde{\mathbf{x}})$. Then, we compute the MSE residual in latent space between latent vector \mathbf{z} from the original patch and the latent vector $\tilde{\mathbf{z}}$ from the reconstructed patch. This "latent loss" term was added to *izif* to formulate the *izizf* loss as

$$L_{izizf} = \frac{1}{n} \|\mathbf{x} - \tilde{\mathbf{x}}\|^2 + \frac{\gamma}{n_d} \|\mathbf{D}_f - \tilde{\mathbf{D}}_f\|^2 + \frac{1}{d} \|\mathbf{z} - \tilde{\mathbf{z}}\|^2. \quad (4)$$

Training was performed with ResNet-50 and DenseNet-169 encoders for each loss function over 30 epochs. Adam optimizer with $\beta_1 = 0.9$, $\beta_2 = 0.999$ and base learning rate of 0.0001 was used. Batch size of 64 was set. We set $\gamma = 0.1$ based on empirical trials.

E. Model Validation and Testing

For any query patch \mathbf{x}_q and its latent vector \mathbf{z}_q , anomaly scores A_s were computed according to the loss function used for model training. For *izizf* loss, we computed mean absolute deviation (MAD) between original and reconstructed patches in image space and latent space as

$$A_s = \frac{1}{n} \|\mathbf{x}_q - \tilde{\mathbf{x}}_q\| + \frac{1}{d} \|\mathbf{z}_q - \tilde{\mathbf{z}}_q\|. \quad (5)$$

For *izi*, *ziz*, and *izif*, since model training did not include a latent loss term, the anomaly score simplified as

$$A_s = \frac{1}{n} \|\mathbf{x}_q - \tilde{\mathbf{x}}_q\|. \quad (6)$$

Model validation was performed after every epoch of encoder training using normal and abnormal patches from the NIH data. The receiver operating characteristic (ROC) area under curve (AUC) was computed and the epoch producing the

highest AUC value was saved. Model testing was performed using normal and abnormal patches from the JSRT data and test AUC were recorded. In both validation and testing, optimal cut-off point was determined by maximizing the Youden Index $J = \text{sensitivity} + \text{specificity} - 1$ and associated sensitivity and specificity were recorded [15].

III. RESULTS

The validation AUC, sensitivity, and specificity are recorded in Table 1. Training and evaluating with *izizf* loss and anomaly scoring function yielded best validation results for ResNet-50 and DenseNet-169 encoders. Models with ResNet-50 and DenseNet-169 encoders were trained using *izizf* loss. Higher testing AUC was noted for the model with ResNet-50 (AUC: 87.89%; sensitivity: 88.65%; specificity: 73.69%) in comparison to DenseNet-169 (AUC: 85.24%; sensitivity: 87.94%; specificity: 69.94%). ROC curves for testing data are shown in Fig. 2a. Sample visual inspection of normal and abnormal testing patches are shown in Fig. 2b and 2c, respectively. Typical reconstructions for abnormal query images were poorer than for normal. Average processing time per patch was 1.77 ms.

IV. DISCUSSIONS

We developed an unsupervised anomaly detection model for lung nodules in CXR. This model was trained with only unlabeled healthy data instances. Using external validation and testing data, we demonstrate that our method generalizes well to independent datasets and can robustly discriminate healthy lung patches from those containing lung nodules.

P-AnoGAN introduced substantial modifications to the f-AnoGAN framework. CXR patches contain diverse anatomy present at multiple scales (e.g., parenchyma, ribs, chest wall, and vasculature). Compared with previous GAN topologies (DCGAN, WGAN), PGAN enabled progressive learning of large scale features followed by finer details making it suitable for our purpose as suggested in [14]. Furthermore, we proposed a new loss function and anomaly scoring function (*izizf*) to jointly quantify deviations between original and reconstructed images in image space, latent space, and discriminator features which improved model performance. Very recently, Nakao et al. used a progressive growing auto-encoding GAN, a CNN encoder, and a "code discriminator" which differentiated randomly sampled latent vectors and latent obtained from real images to detect a variety of abnormalities (including lung masses) in CXR. However,

TABLE I
MODEL PERFORMANCE FOR VALIDATION DATA

Encoder	Experiments			
	Loss	AUC	Sensitivity	Specificity
ResNet-50	<i>ziz</i>	88.58%	93.15%	74.13%
	<i>izi</i>	86.60%	83.56%	74.79%
	<i>izif</i>	88.14%	87.67%	77.50%
	<i>izizf</i>	91.17%	89.04%	80.92%
DenseNet-169	<i>ziz</i>	88.61%	90.41%	77.29%
	<i>izi</i>	86.02%	75.34%	80.24%
	<i>izif</i>	86.48%	84.93%	74.47%
	<i>izizf</i>	91.10%	91.78%	80.60%

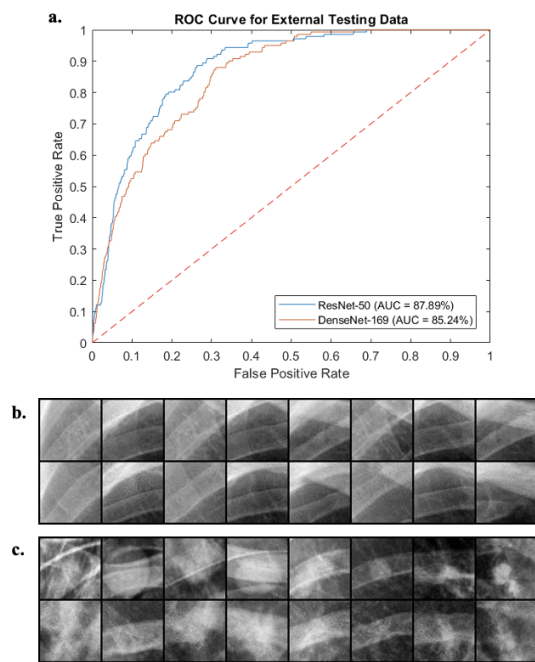


Fig. 2. Summary of test results. (a) ROC curves for models trained using $iziz_f$ loss. Original and reconstructed patches using model trained with ResNet-50 and $iziz_f$ loss for (b) normal and (c) abnormal test patches.

the authors do not provide quantitative results specifically for lung masses. Moreover, their model performance was significantly poorer for subtle pathology (AUC= 70.4%) vs. large pneumonia opacities (AUC= 75.2%). Furthermore, neither of the above-mentioned studies have performed external validation/testing. In contrast, we demonstrated strong performance on the challenging JSRT dataset (avg. manual expert AUC= 84.9%) [11] which indicates robust and generalizable detection of subtle lung nodules in external datasets.

We noted that ResNet-50 yielded higher AUC in both external validation and test data when compared to DenseNet-169. Furthermore, the addition of latent loss in $iziz_f$ loss function and anomaly scoring improved performance. During encoder training, enforcing similarity between latent vectors derived from original and reconstructed patches allowed a robust mapping between image and latent space. Moreover, since encoder training was performed using only healthy data, the image-to-latent mapping is accurate only for healthy queries during evaluation; thus, allowing us to better quantify deviations between original and reconstructed unhealthy data instances in both image and latent spaces. Our paper has several limitations. Firstly, we only used frontal view CXR which may cause difficulties in identifying nodules, particularly those found in the hilar region and near the clavicles [3]. Furthermore, by performing lung segmentation in CXR, we excluded some nodules which may appear outside of the lung region when considering frontal view only; this occurs most commonly when nodules are obscured by surrounding anatomy such as the heart. These challenges could be simultaneously overcome by incorporating lateral view CXR into the algorithm design, training, and testing. Our method detects *any anomaly* appearing in radiographs. Therefore, presence of pathologies appearing alongside lung

cancer (e.g., edema, atelectasis, pleural effusion) and medical support devices may also be detected. While we evaluate our method for the task of lung nodule detection, P-AnoGAN could conceivably be extended for detection of multiple pathologies; though this requires further investigation.

Future works should aim to improve detection performance by optimizing hyper-parameters as well as experiment with additional GANs (StyleGAN) and encoders (InceptionNet, SqueezeNet). We anticipate that such unsupervised deep models can potentially be appealing for triaging and diagnostic tasks in high-volume clinical settings.

ACKNOWLEDGMENTS – The authors thank Dr. Ana M. Ramón Prados for confirming healthy patient cases included in validation data and Dr. Thomas Willett for feedback in the preparation of this work.

REFERENCES

- [1] F. Bray, J. Ferlay, I. Soerjomataram, R. L. Siegel, L. A. Torre, and A. Jemal, "Global cancer statistics 2018: GLOBOCAN estimates of incidence and mortality worldwide for 36 cancers in 185 countries," *CA. Cancer J. Clin.*, vol. 68, no. 6, pp. 394–424, Nov. 2018.
- [2] NLST team, "Results of Initial Low-Dose Computed Tomographic Screening for Lung Cancer," *N. Engl. J. Med.*, vol. 368, no. 21, pp. 1980–1991, May 2013.
- [3] A. del Ciello, P. Franchi, A. Contegiacomo, G. Cicchetti, L. Bonomo, and A. R. Larici, "Missed lung cancer: when, where, and why?," *Diagnostic Interv. Radiol.*, vol. 23, no. 2, pp. 118–126, Mar. 2017.
- [4] Y. Sim et al., "Deep Convolutional Neural Network-based Software Improves Radiologist Detection of Malignant Lung Nodules on Chest Radiographs," *Radiology*, vol. 294, no. 1, pp. 199–209, Jan. 2020.
- [5] J. G. Nam et al., "Development and Validation of Deep Learning-based Automatic Detection Algorithm for Malignant Pulmonary Nodules on Chest Radiographs," *Radiology*, vol. 290, no. 1, pp. 218–228, Jan. 2019.
- [6] T. Schlegl, P. Seeböck, S. M. Waldstein, G. Langs, and U. Schmidt-Erfurth, "f-AnoGAN: Fast unsupervised anomaly detection with generative adversarial networks," *Med. Image Anal.*, vol. 54, pp. 30–44, May 2019.
- [7] L. Sun, J. Wang, Y. Huang, X. Ding, H. Greenspan, and J. Paisley, "An Adversarial Learning Approach to Medical Image Synthesis for Lesion Detection," *IEEE J. Biomed. Heal. Informatics*, vol. 24, no. 8, pp. 2303–2314, Aug. 2020.
- [8] T. Nakao et al., "Unsupervised Deep Anomaly Detection in Chest Radiographs," *J. Digit. Imaging*, Feb. 2021.
- [9] D. Demner-Fushman et al., "Preparing a collection of radiology examinations for distribution and retrieval," *J. Am. Med. Informatics Assoc.*, vol. 23, no. 2, pp. 304–310, Mar. 2016.
- [10] X. Wang, Y. Peng, L. Lu, Z. Lu, M. Bagheri, and R. M. Summers, "ChestX-ray: Hospital-Scale Chest X-ray Database and Benchmarks on Weakly Supervised Classification and Localization of Common Thorax Diseases," *Adv. Comput. Vis. Pattern Recognit.*, pp. 369–392, 2019.
- [11] J. Shiraishi et al., "Development of a Digital Image Database for Chest Radiographs With and Without a Lung Nodule," *Am. J. Roentgenol.*, vol. 174, no. 1, pp. 71–74, Jan. 2000.
- [12] J. A. Stark, "Adaptive image contrast enhancement using generalizations of histogram equalization," *IEEE Trans. Image Process.*, vol. 9, no. 5, pp. 889–896, May 2000.
- [13] Y.-B. Tang, Y.-X. Tang, J. Xiao, and R. M. Summers, "XLSor: A Robust and Accurate Lung Segmentor on Chest X-Rays Using Criss-Cross Attention and Customized Radiorealistic Abnormalities Generation," in *Proc. of 2nd International Conference on Medical Imaging with Deep Learning*, 2019, pp. 457–467.
- [14] T. Karras, T. Aila, S. Laine, and J. Lehtinen, "Progressive growing of GANs for improved quality, stability, and variation," in *6th Int. Conf. Learn. Represent. ICLR 2018*, pp. 1–26, 2018.
- [15] R. Fluss, D. Faraggi, and B. Reiser, "Estimation of the Youden Index and its Associated Cutoff Point," *Biometrical J.*, vol. 47, no. 4, pp. 458–472, Aug. 2005.