



# Utilizing Synthetic Nodules for Improving Nodule Detection in Chest Radiographs

Minki Chung<sup>1</sup> · Seo Taek Kong<sup>1</sup> · Beomhee Park<sup>1</sup> · Younjoon Chung<sup>1</sup> · Kyu-Hwan Jung<sup>1</sup>  · Joon Beom Seo<sup>2</sup>

Received: 26 May 2021 / Revised: 31 January 2022 / Accepted: 14 February 2022 / Published online: 18 March 2022  
© The Author(s) under exclusive licence to Society for Imaging Informatics in Medicine 2022

## Abstract

Algorithms that automatically identify nodular patterns in chest X-ray (CXR) images could benefit radiologists by reducing reading time and improving accuracy. A promising approach is to use deep learning, where a deep neural network (DNN) is trained to classify and localize nodular patterns (including mass) in CXR images. Such algorithms, however, require enough abnormal cases to learn representations of nodular patterns arising in practical clinical settings. Obtaining large amounts of high-quality data is impractical in medical imaging where (1) acquiring labeled images is extremely expensive, (2) annotations are subject to inaccuracies due to the inherent difficulty in interpreting images, and (3) normal cases occur far more frequently than abnormal cases. In this work, we devise a framework to generate realistic nodules and demonstrate how they can be used to train a DNN identify and localize nodular patterns in CXR images. While most previous research applying generative models to medical imaging are limited to generating visually plausible abnormalities and using these patterns for augmentation, we go a step further to show how the training algorithm can be adjusted accordingly to maximally benefit from synthetic abnormal patterns. A high-precision detection model was first developed and tested on internal and external datasets, and the proposed method was shown to enhance the model's recall while retaining the low level of false positives.

**Keywords** Generative adversarial networks · Online data augmentation · Chest radiographs · Computer-aided detection

## Introduction

Chest X-rays (CXRs) are widely used to screen and diagnose pulmonary abnormalities. Among the various abnormalities that can be identified from CXRs, nodules appear as increased attenuation (opacification) patterns and are difficult to interpret, causing high reader variation [1]. With the success of deep learning (DL), interest has surged in developing algorithms to detect nodular patterns from CXRs [2, 3]. However, annotated radiographs are hard to collect: only expert radiologists can label the data, and deep neural networks (DNNs) that require large amounts of labeled data

are either expensive to develop or, having been trained on a small training set, perform subpar to human experts.

Strong and large quantities of appropriate augmentation schemes directly translate to better performance as long as the underlying classes are not perturbed [4, 5]. To preserve objects present in an image, most augmentations transform or distort the background and control the level of noise to the extent that contents are unaffected. In classifying natural images, more sophisticated techniques involving the use of generative algorithms [6] have been devised to augment the training set with synthetic samples [7]. When there is a shortage of certain classes, however, these techniques do not suffice, because of either a lack of diversity of the particular object or class-imbalance. Both issues are prevalent in medical imaging, but additional domain-specific knowledge authorizes the use of targeted augmentation.

In this work, we build upon the above motivations and devise a training scheme to augment the dataset to better localize nodular patterns in CXRs. By incorporating the knowledge that nodule lesions manifest as increased opacity patterns with sizes less than 3cm and are positioned in the interior of lung parenchyma, we develop a synthetic


---

Minki Chung and Seo Taek Kong contributed equally to this work.

✉ Kyu-Hwan Jung  
khwan.jung@vuno.co

<sup>1</sup> VUNO, Seoul, Republic of Korea

<sup>2</sup> Department of Radiology and Research Institute of Radiology, University of Ulsan College of Medicine, Asan Medical Center, Seoul, Korea



**Fig. 1** Examples of synthetically generated nodular patterns: (Top) normal image templates, (middle) masks extracted from real nodular patterns, and (bottom) synthetically generated nodular cases

nodular generation algorithm to augment the training set online. While this work is not the first to apply generative algorithms to medical imaging [8–15], prior works focus on the visual plausibility of abnormalities and have not gone the extra step to investigate how training schemes can be modified accordingly. Our augmentation algorithm is developed without sacrificing the visual plausibility of synthetically generated patterns (see Fig. 1), and we systematically analyze how synthetic nodules can be used to maximize detection performance.

We first developed a standard nodule detection model that achieves high recall at false positives per image (FPPI) comparable to general radiologists when tested on the Japanese Society of Radiological Technology (JSRT) [16] dataset. Other learning-based algorithms yield 25× more FPPI when tested on the same dataset. Then, we show that the augmentation scheme consistently enhances the performance of a detection network across different dataset sizes, as well as levels of class-imbalance within each batch. The latter observation sheds light into how to perform batch sampling for datasets with high class-imbalance. The augmentation scheme developed in this work is then shown to improve recall while maintaining the same level of false positives.

## Materials and Method

### Chest X-ray Images

The institutional review board for human investigation at the **anonymized** hospital approved the study protocol, removed all patient identifiers, and waived informed consent requirements because of the retrospective design of this study. Posteroanterior chest X-rays identified with pulmonary nodules and without lesions were retrospectively collected between January 2010 and November 2016 at **anonymized** hospitals for a total of 1958 nodular and 16,531 non-nodular images. The dataset was split to obtain an internal validation set with 175 nodular and 2065 non-nodular images without patient overlap with the training set. One-hundred fifty-four nodular and 93 normal high-resolution, digitized film CXR images in the Japanese Society of Radiological Technology (JSRT) dataset [16] were used for external validation. Images in this dataset were collected across 13 medical centers in Japan and one institution in the USA, and every abnormal image contains one nodule object. All images were resized to  $1024 \times 1024$  using bi-linear interpolation unless stated otherwise, and normalized after applying windowing to remove irrelevant, outlier pixels.

## Nodule Synthesizer

### Overview

Nodular patterns can be diverse, and the goal is to present a real-time augmentation strategy to generate diverse yet realistic patterns across regions where nodules occur. This section presents our development of a nodule synthesis model that can produce synthetic images in real time, and describes how the network is used to augment the training set online while training, i.e., without inferring and saving synthetic images prior to training.

A lung segmentation model  $\mathcal{S}$  was first trained to segment the lung parenchyma in CXR images on a linear combination of cross-entropy and dice losses. Its predictions  $s_N = \mathcal{S}(I_N)$  on normal images were eroded  $s_N \mapsto \tilde{s}_N$  to serve as background templates on which the generator’s outputs lie. We denote a crop centered at a point in  $p \in \{p : \tilde{s}_N(p) = 1\}$  as  $x$ , and sample a mask  $M \sim \Psi$  from all nodule masks  $\Psi$  corresponding to abnormal images in the training set. The crop is masked  $x \mapsto x_M$  which is then processed through the nodule generator to obtain a synthesized patch  $\mathcal{G}(x_M)$  before reverting to a synthesized image. The overall procedure of generating nodular patterns from normal images and real nodular masks is pictorially summarized in Fig. 2.

### Generator Network

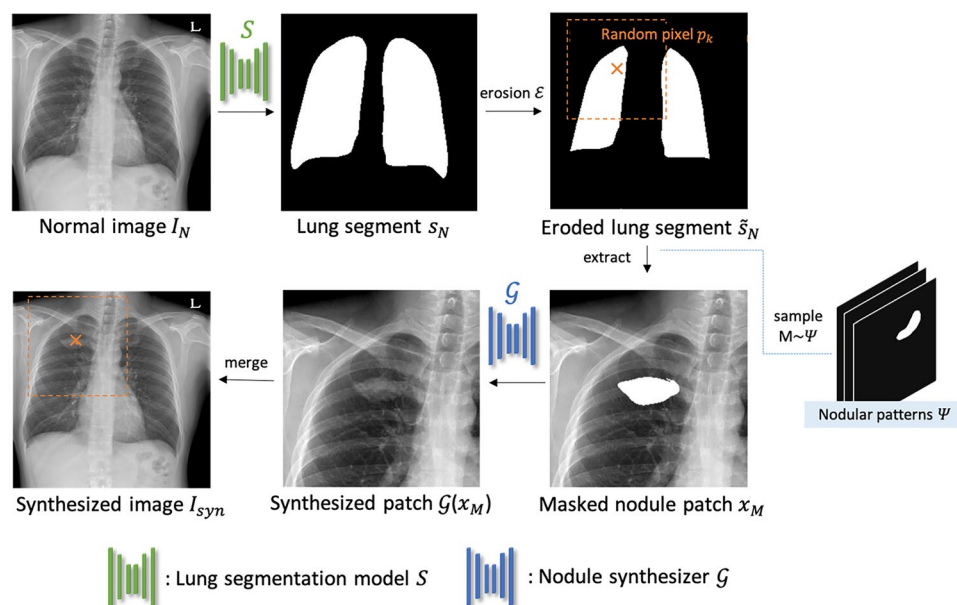
The nodule generator architecture operates in two steps, a coarse synthesizer with a ResNet-34 backbone followed by a refine synthesizer. Synthesizing nodular patterns in normal CXR images is extremely difficult because fine

details may be key factors in distinguishing nodular patterns from morphological features in the lung. Nodular patterns including mass in CXR images might be non-local, where features may have large spatial dependencies. Therefore, a generator must be able to capture distant patterns, but traditional convolution layers encode only local dependencies and do not suffice. Furthermore, the masking operation  $x \mapsto x_M$  loses information, for example vein and bone locations.

The former issue was resolved using contextual attention modules known to capture long-range spatial dependencies dedicated to refine coarse outputs [17]. Contextual attention is computed over the input feature and its masked location as the foreground object, and can therefore incorporate global pathological features in generating nodular patterns. This has been used when capturing non-local background features is pivotal in generating fine-detailed contents. We handled the latter issue by replacing all convolutional layers with gated convolution [18]. Gated convolution mitigates uniform weighting due to binary masks by adaptively updating masks with soft targets and was originally proposed to handle random occlusions. This module was used to realize arbitrary nodular patterns. The resulting nodule synthesizer  $\mathcal{G} : M \circ x \mapsto x'$  training diagram is visualized in Fig. 3.

Lung segments  $s_N$  are extracted from a normal image and eroded, denoted as  $\tilde{s}_N$ . A nodule synthesizer was then trained to generate synthetic nodular patterns  $\mathcal{G}(x_M)$  on the desired position indicated by a mask  $M$  on the LS-GAN loss [19]. The nodule synthesizer architecture consists of an encoder and two decoders, with feature encoding at convolutional layers of ResNet-34 and decoder being nearest-neighbor up-sampling layers for fast inference.

**Fig. 2** Nodular pattern generation schematic diagram. After eroding the lung templates extracted from a lung segmentation model  $\mathcal{S}$ , the nodule synthesizer  $\mathcal{G}$  produces artificial nodular patterns on a masked input  $x_M$ . Masks retrieved from real nodular patterns are randomly placed to fit the lung template



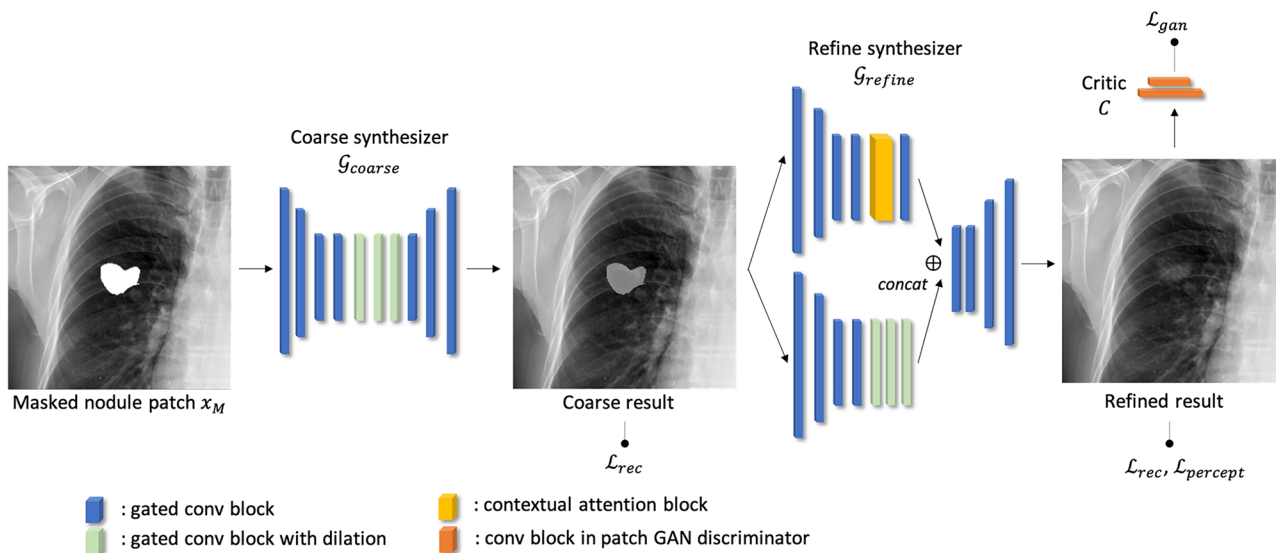


Fig. 3 Nodule synthesizer network  $\mathcal{G}$  and its training procedure

**Online Nodule Augmentation**

Synthetic nodular patterns are to be generated to augment the dataset while training. Offline generation of synthetic patterns limits the diversity of patterns and locations of synthetic nodular patterns because the number of synthetic nodules that help train the detection network must be known in advance. This reduces the flexibility of modifying sampling ratios among normal or real and synthetic nodules. An overview of the algorithm is presented in Algorithm 1.

At every SGD step, a batch  $B' = \{I_n, I_a b\}$  of normal  $I_n$  and abnormal images  $I_a b$  are sampled and synthetic nodules are generated on top of background images within the batch to obtain an augmented batch  $B_{s,yn} = \{I_n, I_a b, I_{s,yn}\}$ . The number of objects  $K \sim Geom(1/\mu)$  to be generated on a template  $I_n \in B'$  is first randomly sampled for each candidate synthetic image. Masks  $M = \{M_k\}_{k=1}^K$  are then drawn from the set of real nodular masks  $\Psi$  uniformly at random with each mask centered at  $p_k$  randomly drawn within the lung such that masks do not overlap. Lung regions  $\mathcal{S} : I_n \mapsto \mathcal{S}$

**Algorithm 1** Online Nodule Augmentation

**Inputs:** Lung segmentation model  $\mathcal{S}$ , nodule synthesizer  $\mathcal{G}$ , detection network  $\mathcal{F}$ , abnormal dataset  $\mathcal{D}_a$ , normal dataset  $\mathcal{D}_n$ , set of feasible masks  $\Psi$ .

**repeat**

$\mathcal{B}_a = \{(I_a^{(i)}, y_a^{(i)})\}_{i=1}^{m_a} \sim \mathcal{D}_a$ .

$\mathcal{B}_n = \{(I_n^{(i)}, y_n^{(i)})\}_{i=1}^{m_n} \sim \mathcal{D}_n$ .

$x_n^{(i)} \leftarrow \mathcal{S}(I_n^{(i)})$

**for**  $i = 1, \dots, m_s$  **do**

$K \sim Geom(1/\mu)$  ▷ Number of synthetic objects.

$M_1, \dots, M_K \sim Uniform(\Psi)$  ▷ Sample nodular patterns.

**for**  $k = 1, \dots, K$  **do**

$x_n^{(i)} \leftarrow \mathcal{E}(x_n^{(i)}, |M_k|)$  ▷ Erode the lung interior.

**end for**

$p_1, \dots, p_K \sim Uniform(x_n)$  ▷ Sample object centers.

$y^{(i)} \leftarrow \sum_{k=1}^K M_k(p_k)$  ▷ Aggregate masks.

$I_s^{(i)} \leftarrow \mathcal{G}(y^{(i)} \circ x_n^{(i)})$  ▷ Generate synthetic image.

**end for**

$\mathcal{B}_s \leftarrow \{(I_s^{(i)}, y^{(i)})\}_{i=1}^{m_s}$

$\min_{\theta} \mathcal{L}(\theta; \mathcal{B}_a, \mathcal{B}_n, \mathcal{B}_s)$  ▷ Update model on the augmented batch.

**until** convergence criterion is met

extracted from a reference template  $I_n$  (image) are eroded  $\mathcal{E} : (\tilde{s}, |M_k|) \mapsto x_M$  recursively using the mask's size  $|M_k|$  to obtain the masked nodule patch  $x_M$  with objects  $y = \sum_k M_k$ . The nodule synthesizer then generates the synthetic image  $(\tilde{s}, x_M) \mapsto \mathcal{G}(x_M)$ . Together, synthetic patterns are guaranteed to be in actual locations where nodular patterns occur and are diverse enough to help the detection network generalize unseen patterns.

## Experiments

### Implementation Details

RetinaNet [20] with a ResNet-50 backbone detection network was trained on the focal loss to localize nodular (abnormal) lesions in patches. When training both standard and synthetic augmentation models, rotation  $\pm 10$ , (0.9, 1.1), blurring, and sharpening augmentations were applied. Nodule synthesizer and critic networks were trained using alternating descent with the Adam optimizer with learning rates  $10^{-4}$  and  $4 \cdot 10^{-4}$ , respectively, decayed by a factor of 2 every 10 epochs.

When training the synthesizer, we rescaled the images by a factor of 2 on both width and height, applying zero-padding whenever the patterns' centers were located near image boundaries. The training objective for each sample  $X$  was a combination of reconstruction, perceptual, and adversarial losses to solve

$$\mathcal{G} = \arg \min_G \max_{D: \mathcal{I} \rightarrow [0,1]} \lambda_r \|G(M \circ I) - X\|_2 + \lambda_p \|\phi(G(M \circ I)) - \phi(I)\|_2^2 + \lambda_a \mathbb{E}_{I \sim \mathbb{P}_{real}} [D(G(M \circ I)) + (1 - D(I))]$$

with coefficients  $\lambda_r = 1$ ,  $\lambda_p = 10$ , and  $\lambda_a = 0.01$ , respectively. Here,  $\phi$  is a normalized feature extraction network

as in [21–23]. This discriminative network was trained using the Adam optimizer with learning rate  $10^{-5}$  to minimize focal and  $L_1$  loss for the classification and regression heads respectively. The mean number of objects in synthetic images was set as  $\mu = 1.25$ .

## Results

Here, we compare the detection model we developed without and with synthetic augmentation to several data-driven algorithms that report their performance on the JSRT dataset in Table 1. It is difficult to compare only the performance enhancements due to algorithmic differences because baselines were developed on different train/validation sets and a few were tested only on a subset of the JSRT dataset, but we list them for reference. General and chest radiologists reportedly achieved 64% and 77% recall at 0.072 and 0.076 false positives per image, respectively [16]. In comparison, other baselines [24–29] we found report 25× more FPPI than human experts, whereas both our detection models can attain 0.08 FPPI. This low false positive level is necessary for practical deployment of computer-aided diagnostic systems [1, 30]. At this level of false positives, the detection model was observed to detect nodular patterns at 49% and 52% recall with and without synthetic augmentation. This enhancement in recall shows that synthetic augmentation can improve recall while retaining the level of false positives.

### Dataset Size


To understand the relationship between dataset size and how synthesized nodules affect training a detection model, we fixed the normal:abnormal:synthetic data ratio within a batch to 2 : 1 : 1 and varied the training set size  $N$  used to train both the generative and detection models. The

**Table 1** Recall and false positives per image reported on the JSRT dataset, listed in order of increasing false positive rates. All data-driven methods other than that developed in this work exceed false positives per image by a factor of  $\geq 25$  of our model and radiologists.

Methods	FP/image	Recall (%)	Training database
General radiologists [16]	0.072	64	-
Chest radiologists [16]	0.076	77	-
<b>Standard (ours)</b>	0.08	49.4	In-house (w/o JSRT)
<b>Synthetic Aug. (ours)</b>	0.08	52.0	In-house
Schiham et al. [28]	2.0	51	Nodular JSRT
Hardie et al. [26]	2.0	63	140 nodular JSRT
Chen et al. [24]	2.0	64.9	140 normal and nodular JSRT
Li et al. (Single model) [27]	2.1	57	Full JSRT
Coppini et al. [25]	4.3	60	Nodular JSRT
Wei et al. [29]	5.4	80	Full JSRT

Note that our work and others' used different training datasets and evaluation schemes (e.g.,  $K$ -fold validation) and should be taken only as references as they are not directly comparable

**Fig. 4** Synthesized images when the nodule synthesis network was trained on (left) 250, (middle) 500, and (right) 1958 abnormal images. Synthetic examples become more visually realistic when the size of training image is increased



**Table 2** Average precision improvements as training dataset size varies. Synthetic generation of abnormal data consistently enhances the detection network's precision in both low-data and large-data regimes

$n, ab$	Internal validation	JSRT
16, 531, 1958	75.1 → 77.8	47.0 → 52.3
2500, 500	64.3 → 70.8	39.8 → 47.8
1250, 250	60.5 → 64.3	32.9 → 36.3

generative capability of the nodule synthesizer is shown in Fig. 4 visualizing how the generator becomes more capable of generating realistic synthesized nodules as the number of abnormal training samples increases. The detection model's average precision (AP) change when using standard → synthetic augmentation on an internal and JSRT validation set are shown in Table 2, where it is shown that synthetic augmentation always helps generalize throughout various dataset sizes and internal/external sets.

### Batch Sampling Ratio

While there are no restrictions on the class of nodular patterns our generator can produce, the number of real data (template and nodular patterns) available in the training set will always remain fixed. To observe the extent to which a discriminative network can benefit from using synthetic

**Table 3** Ablation study: performances when trained using different sampling ratios within a batch. Note that recall and precision performances are at 0.08 FPPI. Batches containing synthetic samples enhances performances while batches with only synthetic samples shows drastic performance drop. Different batch ratio strategy should be taken depending on the target metric. Bold numbers denotes the highest performance in each performance metric

$n : ab : syn$	AP	Recall	Precision
Standard	75.1	82.3	47.7
3 : 2 : 1	75.4	82.3	47.7
2 : 1 : 1	<b>77.8</b>	83.3	48.0
3 : 1 : 2	75.2	<b>84.3</b>	<b>48.3</b>
1 : 0 : 1	23.9	29.8	24.8

nodules, we carefully increased the ratio between the real ( $ab$ ) to synthetic ( $syn$ ) abnormal images while keeping the number of normal ( $n$ ) images equal to their sum. The model's performances on the in-house dataset are shown in Table 3. We noticed that a DNN trained using the smallest ratio of synthetic images already outperforms the standard model trained using only real images, but the performance starts to degrade as the number of synthetic nodules excessively exceeds the real abnormal samples. This concave performance indicates how a small number of synthetic patterns are helpful, but not particularly effective, and as more synthetic images are merged with the real dataset, it hinders training. The relatively low precision on the internal dataset is due to non-nodular lesions including focal interstitial opacity or consolidation considered normal for our nodule detection task.

We evaluate the performance of a model trained using only on real normal and synthetic abnormal images to show a limiting case in Table 3. While synthetic images can aid in learning discriminative features, this experiment shows that synthetic samples could not be used as a replacement to real abnormalities. This is consistent with the observation that a generative model learns a compressed representation of the real dataset.

### Conclusion

We proposed a generative framework which can be used for real-time augmentation to generate synthetic nodules. With the use of synthetic nodular patterns, batch sampling can be performed with nearly even classes which was shown to enhance the recall of a highly precise pulmonary nodule detection model. Our experiments illustrate the factors that must be considered in training a detection network to localize nodules using the synthesized images and showed the effectiveness of our approach on internal and external datasets. By controlling various dataset sizes to train the generative and detection networks, we showed the trade-off between the generative capacity and detection performance enhancements resulting from synthesized nodules. We hope

that our experiments can guide future research in using synthetic images to aid multi-label classification and localization in CXR images which are known to suffer severely from class imbalance.

**Funding** Not applicable.

**Availability of Data and Material** Not applicable.

**Code Availability** Not applicable.

## Declarations

**Conflict of Interest** Minki Chung, Seo Taek Kong, Beomhee Park, Younjoon Chung and Kyu-Hwan Jung are employees of VUNO Inc. Kyu-Hwan Jung is an equity holder of VUNO Inc.

## References

- Sung, J., Park, S., Lee, S.M., Bae, W., Park, B., Jung, E., Seo, J.B., Jung, K.H.: Added value of deep learning–based detection system for multiple major findings on chest radiographs: A randomized crossover study. *Radiology* **299**(2), 450–459 (2021)
- Irvin, J., Rajpurkar, P., Ko, M., Yu, Y., Ciurea-Ilicus, S., Chute, C., Marklund, H., Haghighi, B., Ball, R., Shpankaya, K., et al.: Chexpert: A large chest radiograph dataset with uncertainty labels and expert comparison. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 33, pp. 590–597 (2019)
- Wang, X., Peng, Y., Lu, L., Lu, Z., Bagheri, M., Summers, R.M.: Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 2097–2106 (2017)
- Berthelot, D., Carlini, N., Cubuk, E.D., Kurakin, A., Sohn, K., Zhang, H., Raffel, C.: Remixmatch: Semi-supervised learning with distribution matching and augmentation anchoring. In: *International Conference on Learning Representations* (2020)
- Cubuk, E.D., Zoph, B., Shlens, J., Le, Q.V.: Randaugment: Practical automated data augmentation with a reduced search space. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*. pp. 702–703 (2020)
- Goodfellow, I.J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2*. p. 2672–2680. NIPS'14, MIT Press, Cambridge, MA, USA (2014)
- Shrivastava, A., Pfister, T., Tuzel, O., Susskind, J., Wang, W., Webb, R.: Learning from simulated and unsupervised images through adversarial training. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 2107–2116 (2017)
- Chuquicusma, M.J., Hussein, S., Burt, J., Bagci, U.: How to fool radiologists with generative adversarial networks? a visual turing test for lung cancer diagnosis. In: *2018 IEEE 15th international symposium on biomedical imaging (ISBI 2018)*. pp. 240–244. IEEE (2018)
- Frid-Adar, M., Amer, R., Greenspan, H.: Endotracheal tube detection and segmentation in chest radiographs using synthetic data. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 784–792. Springer (2019)
- Madani, A., Moradi, M., Karargyris, A., Syeda-Mahmood, T.: Chest x-ray generation and data augmentation for cardiovascular abnormality classification. In: *Medical Imaging 2018: Image Processing*. vol. 10574, p. 105741M. International Society for Optics and Photonics (2018)
- Madani, A., Moradi, M., Karargyris, A., Syeda-Mahmood, T.: Semi-supervised learning with generative adversarial networks for chest x-ray classification with ability of data domain adaptation. In: *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*. pp. 1038–1042. IEEE (2018)
- Nie, D., Trullo, R., Lian, J., Petitjean, C., Ruan, S., Wang, Q., Shen, D.: Medical image synthesis with context-aware generative adversarial networks. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 417–425. Springer (2017)
- Sandfort, V., Yan, K., Pickhardt, P.J., Summers, R.M.: Data augmentation using generative adversarial networks (cycleGAN) to improve generalizability in ct segmentation tasks. *Scientific reports* **9**(1), 1–9 (2019)
- Wu, E., Wu, K., Cox, D., Lotter, W.: Conditional infilling GANs for data augmentation in mammogram classification. In: *Image Analysis for Moving Organ, Breast, and Thoracic Images*, pp. 98–106. Springer (2018)
- Xing, Y., Ge, Z., Zeng, R., Mahapatra, D., Seah, J., Law, M., Drummond, T.: Adversarial pulmonary pathology translation for pairwise chest x-ray data augmentation. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 757–765. Springer (2019)
- Shiraishi, J., Katsuragawa, S., Ikezoe, J., Matsumoto, T., Kobayashi, T., Komatsu, K.I., Matsui, M., Fujita, H., Kodera, Y., Doi, K.: 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. *American Journal of Roentgenology* **174**(1), 71–74 (2000)
- Yu, J., Lin, Z., Yang, J., Shen, X., Lu, X., Huang, T.S.: Generative image inpainting with contextual attention. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 5505–5514 (2018)
- Yu, J., Lin, Z., Yang, J., Shen, X., Lu, X., Huang, T.S.: Free-form image inpainting with gated convolution. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 4471–4480 (2019)
- Mao, X., Li, Q., Xie, H., Lau, R.Y., Wang, Z., Paul Smolley, S.: Least squares generative adversarial networks. In: *Proceedings of the IEEE international conference on computer vision*. pp. 2794–2802 (2017)
- Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: *Proceedings of the IEEE international conference on computer vision*. pp. 2980–2988 (2017)
- Gatys, L.A., Ecker, A.S., Bethge, M.: Image style transfer using convolutional neural networks. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2016)
- Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual losses for real-time style transfer and super-resolution. In: *European conference on computer vision*. pp. 694–711. Springer (2016)
- Ledig, C., Theis, L., Huszár, F., Caballero, J., Cunningham, A., Acosta, A., Aitken, A., Tejani, A., Totz, J., Wang, Z., 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*. pp. 4681–4690 (2017)
- Chen, S., Suzuki, K., MacMahon, H.: Development and evaluation of a computer-aided diagnostic scheme for lung nodule detection in chest radiographs by means of two-stage nodule enhancement with support vector classification. *Medical physics* **38**(4), 1844–1858 (2011)

25. Coppini, G., Diciotti, S., Falchini, M., Villari, N., Valli, G.: Neural networks for computer-aided diagnosis: detection of lung nodules in chest radiograms. *IEEE Transactions on Information Technology in Biomedicine* **7**(4), 344–357 (2003)
26. Hardie, R.C., Rogers, S.K., Wilson, T., Rogers, A.: Performance analysis of a new computer aided detection system for identifying lung nodules on chest radiographs. *Medical Image Analysis* **12**(3), 240–258 (2008)
27. Li, C., Zhu, G., Wu, X., Wang, Y.: False-positive reduction on lung nodules detection in chest radiographs by ensemble of convolutional neural networks. *IEEE Access* **6**, 16060–16067 (2018)
28. Schilham, A.M., Van Ginneken, B., Loog, M.: A computer-aided diagnosis system for detection of lung nodules in chest radiographs with an evaluation on a public database. *Medical Image Analysis* **10**(2), 247–258 (2006)
29. Wei, J., Hagihara, Y., Shimizu, A., Kobatake, H.: Optimal image feature set for detecting lung nodules on chest x-ray images. In: *CARS 2002 computer assisted radiology and surgery*, pp. 706–711. Springer (2002)
30. Park, S., Park, G., Lee, S.M., Kim, W., Park, H., Jung, K., Seo, J.B.: Deep learning-based differentiation of invasive adenocarcinomas from preinvasive or minimally invasive lesions among pulmonary subsolid nodules. *European Radiology* pp. 1–9 (2021)

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.