GENERATION OF ANONYMOUS CHEST RADIOGRAPHS USING LATENT DIFFUSION MODELS FOR TRAINING THORACIC ABNORMALITY CLASSIFICATION SYSTEMS

Kai Packhäuser, Lukas Folle, Florian Thamm, Andreas Maier

Pattern Recognition Lab Department of Computer Science Friedrich-Alexander-Universität Erlangen-Nürnberg

ABSTRACT

The availability of large-scale chest X-ray datasets is a requirement for developing well-performing deep learningbased algorithms in thoracic abnormality detection and classification. However, biometric identifiers in chest radiographs hinder the public sharing of such data for research purposes due to the risk of patient re-identification. To counteract this issue, synthetic data generation offers a solution for anonymizing medical images. This work employs a latent diffusion model to synthesize an anonymous chest X-ray dataset of high-quality class-conditional images. We propose a privacy-enhancing sampling strategy to ensure the non-transference of biometric information during the image generation process. The quality of the generated images and the feasibility of serving as exclusive training data are evaluated on a thoracic abnormality classification task. Compared to a real classifier, we achieve competitive results with a performance gap of only 3.5% in the area under the receiver operating characteristic curve.

Index Terms— Chest Radiography, Synthetic Image Generation, Abnormality Classification, Patient Privacy

1. INTRODUCTION

Over the past years, the availability of novel deep learning (DL) [1] techniques has led to significant advances in the development of diagnostic algorithms for the detection and classification of thoracic abnormalities [2]. As the training of DL models is data-driven, it typically requires enormous amounts of training samples to obtain well-performing networks. However, medical images, e. g., chest radiographs, include personal information and are subject to certain privacy protection regulations, such as the EU General Data Protection Regulation (GDPR), which hinders the widespread sharing and, thus, utilization of acquired data. Therefore, to make large-scale medical image collections publicly available for research purposes, it is necessary to strictly comply with applicable objectives and requirements for the robust and reliable anonymization of sensitive patient data. Conventional techniques for anonymizing chest radiographs rely on removing personally identifiable meta-information, e.g., patient names, or replacing them with pseudonyms. Furthermore, in the image domain itself, critical areas are typically obscured with black boxes. In one of our recent studies [3], however, we empirically demonstrate that such approaches do not offer sufficient means of protecting patient privacy. This is primarily because chest radiographs inherently contain biometric information (similar to a fingerprint) that can be exploited to successfully re-identify specific patients by DL-based linkage attacks. Hence, there is an urgent need for more sophisticated privacy-enhancing techniques (PETs) in medical imaging.

In this context, the generation of synthetic medical images with privacy guarantees has attracted attention as a promising solution for anonymization and overcoming data-sharing limitations. Generative models learn the probability distribution of a given real dataset and can be used to synthesize realistic representations by sampling from the learned data distribution. Provided that synthetic medical images are of sufficient quality, such data could be used to train DL models at scale. Han et al. [4] utilized progressively growing generative adversarial networks (PGGANs) [5] to synthesize chest radiographs and demonstrated the feasibility of GANbased data augmentations. Moreover, they confirm the utility of exclusively synthetic training data in a clinical setting. Recently, diffusion probabilistic models [6] have become state-of-the-art for tasks such as high-resolution image synthesis. Chambon et al. [7] leverage a latent diffusion model (LDM) [8] for the generation of class-conditional chest radiographs. Nevertheless, the authors do not investigate whether the generated images can convey relevant class information during the training process of an abnormality classifier. Furthermore, their dataset is limited to only 2,000 samples.

In this work, we aim to investigate the feasibility of using exclusively synthetic training datasets for learning and recognizing thoracic abnormalities in chest radiographs. Therefore, we utilize an LDM [8] to generate high-quality class-conditional images from a given real data distribution. To synthesize a fully anonymous dataset from the trained generative model without transferring biometric patient identifiers, we propose a rigorous privacy-enhancing sampling strategy that excludes synthetic images, provided a specific patient identity has been learned from the original training data. Finally, to quantify the quality of the generated anonymous dataset and to evaluate its applicability for training a thoracic abnormality identification system, we conduct multiple classification experiments and compare the performances of classifiers trained on real versus synthetic data, respectively. Throughout our work, we leverage a PGGAN [5] as a synthetic image generation baseline.

2. METHODS

2.1. Dataset

We use the large-scale ChestX-ray14 dataset [9] consisting of 112,120 chest radiographs from 30,805 individual patients. Accompanying metadata provides information about the corresponding 14 abnormality labels, including Atelectasis, Cardiomegaly, Consolidation, Edema, Effusion, Emphysema, Fibrosis, Hernia, Infiltration, Mass, Nodule, Pleural Thickening, Pneumonia, and Pneumothorax. Healthy subjects are labeled with an additional class, indicating that none of the above-mentioned abnormalities are present. To simplify the class-conditional image generation task in this work, we exclude all images with multiple abnormality labels. Furthermore, we limit the maximum number of follow-up scans per patient to 5 to prevent the generative models from learning patient-specific patterns from over-represented subjects. Lastly, as the anatomy of young patients may not yet be fully developed, we only consider images from patients older than 21 years. After this data reduction procedure, the remaining 56,352 images are split into a training, validation, and test set by a ratio of 70:10:20. To this end, a patient-wise splitting strategy is applied to avoid potential patient overlap between the subsets while maintaining the overall class distribution.

2.2. Latent diffusion model

For the class-conditional image generation task, we apply the LDM proposed in [8] leveraging pre-trained autoencoders. In this approach, input images are first embedded into a latent space of size $64 \times 64 \times 3$ using the encoder of a vector quantized-variational autoencoder (VQ-VAE) with 32, 64, and 128 channels in each stage. A diffusion model operates within that lower space dominated by lower frequencies and performs 1000 denoising steps with a U-Net (32, 128, and 256 channels). In the final step, the decoder of the autoencoder increases the spatial resolution to 256×256 pixels while introducing higher frequencies. The class-conditional information is incorporated using a trainable lookup table. This is realized by combining the class embeddings with the diffusion process using cross-attention in the bottleneck of the U-Net. We closely followed the implementation by Rombach et al. [8].



Fig. 1. Illustration of the anonymous image generation process using the proposed privacy-enhancing sampling strategy. We apply generative modeling (1) techniques to synthesize chest radiographs. We then use patient retrieval (2) and verification (3) networks to ensure the non-transference of biometric identifiers. The resulting anonymous dataset maintains the size and class distribution of the original training set and is used for training an abnormality classifier (4).

2.3. Progressively growing GAN

As a synthetic image generation baseline, we employ the PGGAN architecture proposed in [5]. PGGANs are based on the idea of starting the training process with low-resolution images before gradually increasing the image resolution by adding layers to the network architecture. This allows the generator to focus on large-scale image structures at the beginning of the training procedure and to learn finer highresolution details in later stages. We use a generator and discriminator that have the same general structure and grow equally over time. Both networks incorporate replicated 3-layer blocks consisting of an upsampling or downsampling layer, a 3×3 convolutional layer, and a leaky rectified linear unit (LReLU) with a slope of 0.2. We adapt the originally proposed PGGAN architecture to produce chest radiographs with a final spatial resolution of 256×256 pixels. The input to the generator is represented by a 256-dimensional latent noise vector. Its last 15 entries consist of the conditional vector, which controls the image class to be generated. For the discriminator, we use a projection-based approach [10] to incorporate conditional information into the network.

2.4. Privacy-enhancing sampling strategy

While synthetic data is often referred to as anonymous per se, generative models may memorize training examples and reproduce specific patient identities during inference [11]. To ensure the non-transference of biometric patterns when generating synthetic chest radiographs, we propose a rigorous privacy-enhancing sampling strategy (see Fig. 1) that leverages state-of-the-art patient matching approaches. The strategy is based on the idea of excluding synthetic images in case patient-specific biometrics have been reproduced in a generated scan. To realize such an image selection process, we first apply a patient retrieval network for each generated chest radiograph to find its top-1 image in terms of patient similarity from the original training set. Then, we perform a 1-to-1 matching step for the recommended image pairs (real vs. synthetic) by applying a patient verification network that yields the probability of whether or not the synthetic image and the recommended real image belong to the same patient. Based on this probability value, we exclude all synthetic images that exceed a threshold of t = 0.5. We employ pre-trained patient retrieval and verification models proposed in previous work [3]. The used networks achieve a top-1 precision of 99.6% and an area under the receiver operating characteristic curve (AUC) of 99.4%, respectively.

2.5. Thoracic abnormality classification

For the downstream abnormality classification task, we employ CheXNet [12], a densely connected convolutional network (DenseNet) [13] consisting of 121 layers. Its input layer receives chest radiographs with a resolution of 224×224 pixels. The final classification layer yields a 14-dimensional output vector indicating the probability of the presence or absence of each abnormality class appearing in the ChestX-ray14 dataset.

3. EXPERIMENTS

All experiments were conducted using PyTorch and Python. The experimental setup included the following three steps.

3.1. Training of the generative models

The first step of our experimental pipeline consisted of training the generative models. For the LDM, we first trained the autoencoder with a combination of the perceptual loss [14] and a patch-based adversarial objective [15] for 100 epochs using a learning rate of $4.5 \cdot 10^{-6}$. Then, the autoencoder was applied to embed the input images into latent representations, which were used as inputs for the underlying diffusion model. The diffusion model was trained with a learning rate of 10^{-6} using the L_1 loss until no further improvement in the validation loss (250 epochs).

The PGGAN architecture was trained using the Wasserstein loss with a gradient penalty coefficient of 10. For optimization, we used the Adam [16] optimizer with a learning rate of 0.001. The training was initiated at an image resolution of 4×4 . We gradually added layers to both the generator and the discriminator to increase the resolution by factors of 2 until an image size of 256×256 was reached. Each stage of the PGGAN was trained until the generator and discriminator loss values stabilized (100 epochs).

Table 1. Comparison of the abnormality identification performance of CheXNet when using either real or exclusively synthetic training sets. For all evaluations, the same real testing data is employed. We report the means and standard deviations of the resulting AUC values after 10 independent training and testing runs. Best synthetic results are shown in bold.

g and testing ran			
Training set	Real	Syn _{PGGAN}	Syn _{LDM}
Atelectasis	81.3 ± 0.8	70.1 ± 1.2	76.2 ± 0.4
Cardiomegaly	92.9 ± 0.6	86.4 ± 1.4	88.6 ± 0.8
Consolidation	74.8 ± 1.0	68.0 ± 2.8	75.5 ± 1.1
Edema	92.8 ± 0.8	84.4 ± 2.8	87.5 ± 1.7
Effusion	90.7 ± 0.4	83.2 ± 0.9	85.9 ± 0.9
Emphysema	88.1 ± 0.8	76.5 ± 1.6	83.9 ± 1.0
Fibrosis	80.8 ± 1.0	69.4 ± 2.9	77.3 ± 0.6
Hernia	93.5 ± 1.5	80.6 ± 3.5	93.7 ± 1.4
Infiltration	68.7 ± 0.2	59.1 ± 0.6	63.4 ± 0.5
Mass	81.0 ± 1.3	67.7 ± 0.9	76.9 ± 0.8
Nodule	71.4 ± 1.0	60.8 ± 1.2	67.8 ± 0.7
Pleural Thickening	75.5 ± 1.2	68.4 ± 1.3	73.4 ± 0.8
Pneumonia	70.8 ± 3.3	61.5 ± 5.1	65.8 ± 2.0
Pneumothorax	79.9 ± 0.7	70.1 ± 1.5	77.0 ± 0.5
Mean	81.6 ± 0.4	71.9 ± 0.8	78.1 ± 0.3

3.2. Privacy-enhancing image generation

After training the generative models, we synthesized two anonymous chest X-ray datasets by leveraging the trained LDM and PGGAN, respectively. For this purpose, we applied the proposed privacy-enhancing sampling strategy to ensure the non-transference of biometric identifiers during the image generation process. Note that the size and class distribution of the real training and validation set was maintained when creating the synthetic datasets. No oversampling of under-represented classes was performed. Both the LDMbased and the GAN-based datasets were used for further experiments as described in the following step.

3.3. Thoracic abnormality classification

To assess the extent to which synthetic data can convey relevant class information during the training process of a thoracic abnormality classifier, we conducted experiments with both a real training set and with exclusive synthetic training sets generated by either the trained PGGAN or the LDM (Syn_{PGGAN} and Syn_{LDM}). We trained the classifiers using the class-wise binary cross-entropy (BCE) loss and stochastic gradient descent (SGD) with a momentum of 0.9 and weight decay of 0.0001. The initial learning rate was set to 0.01 and was divided by a factor of 10 if the loss did not improve after an epoch. Model training was conducted until early stopping with a patience of 3. Each individually trained classifier was evaluated using the same real testing set. For each scenario, we performed 10 independent training and testing runs. The network performances are compared by reporting the means and standard deviations of the resulting AUC values.



Fig. 2. Randomly selected images generated by (a) the trained PGGAN and (b) the LDM, respectively. The positions and areas of induced abnormalities are highlighted with red arrows or circles.

4. RESULTS

4.1. Quantitative results

The results of the conducted thoracic abnormality classification experiments are summarized in Table 1. As can be seen, the overall best performance is achieved when training the classifier with real images from the ChestX-ray14 dataset. In this case, we obtain a mean AUC of 81.6 %. Compared to this reference, the results significantly decrease for the classifier trained with the synthetic dataset generated by the PGGAN. The mean AUC of 71.9 % indicates that the GAN-based images fail to convey relevant class information during the training of the classification system. In contrast, when using training images generated by the LDM, we observe competitive classification results with a mean AUC of 78.1 % on our test set. Thus, the performance is only slightly below the reference, indicating the ability of the LDM-based chest radiographs to adequately preserve class-specific patterns. Interestingly, the performance for pathology classes Consolidation and Hernia is even higher than the reference.

4.2. Qualitative results

Randomly selected chest radiographs generated by the trained PGGAN and the LDM are shown in Fig. 2. In direct comparison, we perceive better image quality for the LDM-based chest radiographs. The trained LDM can produce high-resolution images with realistic patient anatomy. Moreover, as seen in the figure, each of the presented abnormalities

becomes visible in the generated images. In contrast, the GAN-based chest radiographs show a slightly worse image quality. This becomes apparent, for instance, by the partially unrealistic shape of the ribs and lungs, or through isolated artifacts that appear in the images. Moreover, we observe that the induced class conditions are a bit less pronounced in the images generated by the PGGAN.

5. DISCUSSION AND CONCLUSION

In this paper, we proposed a privacy-enhancing image sampling strategy to synthesize anonymous chest X-ray datasets. We leveraged an LDM and a PGGAN to generate high-quality class-conditional images. The obtained quantitative and qualitative results indicate the competitiveness of diffusion-based chest radiographs for training a thoracic abnormality classification system. With our work, we do not want to suggest that synthetic data can replace real data. In a real-life scenario, this would only be applicable if the quality of synthetic data is on par with real data. However, if this can be guaranteed, we hypothesize that our proposed approach allows for significant mitigation of data-sharing limitations.

In future work, next to further improving the overall image quality, we aim to exploit the generative models for compensating under-represented classes or creating infinitely large synthetic datasets. Lastly, we plan to incorporate a privacy-enhancing mechanism directly into the LDM to ensure the non-transference of biometric information during the training process.

6. COMPLIANCE WITH ETHICAL STANDARDS

This research study was conducted retrospectively using human subject data made available in open access by the National Institutes of Health (NIH) Clinical Center [9]. Ethical approval was not required as confirmed by the license attached with the open-access data.

7. ACKNOWLEDGMENTS

The research leading to these results has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (ERC Grant no. 810316). The authors gratefully acknowledge the scientific support and HPC resources provided by the Erlangen National High Performance Computing Center (NHR@FAU) of the Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU). The hardware is funded by the German Research Foundation (DFG). The authors declare that they have no conflicts of interest.

8. REFERENCES

- Andreas Maier, Christopher Syben, Tobias Lasser, and Christian Riess, "A gentle introduction to deep learning in medical image processing," *Zeitschrift für Medizinische Physik*, vol. 29, no. 2, pp. 86–101, 2019.
- [2] Sebastian Gündel, Sasa Grbic, Bogdan Georgescu, Siqi Liu, Andreas Maier, and Dorin Comaniciu, "Learning to Recognize Abnormalities in Chest X-Rays with Location-Aware Dense Networks," in *Iberoamerican Congress on Pattern Recognition*. Springer, 2018, pp. 757–765.
- [3] Kai Packhäuser, Sebastian Gündel, Nicolas Münster, Christopher Syben, Vincent Christlein, and Andreas Maier, "Deep learning-based patient re-identification is able to exploit the biometric nature of medical chest Xray data," *Scientific Reports*, vol. 12, no. 1, pp. 1–13, 2022.
- [4] Tianyu Han et al., "Breaking medical data sharing boundaries by using synthesized radiographs," *Science Advances*, vol. 6, no. 49, pp. eabb7973, 2020.
- [5] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen, "Progressive Growing of GANs for Improved Quality, Stability, and Variation," in *International Conference on Learning Representations*, 2018.
- [6] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli, "Deep Unsupervised Learning using Nonequilibrium Thermodynamics," in

International Conference on Machine Learning. PMLR, 2015, pp. 2256–2265.

- [7] Pierre Chambon, Christian Bluethgen, Curtis P Langlotz, and Akshay Chaudhari, "Adapting Pretrained Vision-Language Foundational Models to Medical Imaging Domains," *arXiv preprint arXiv:2210.04133*, 2022.
- [8] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer, "High-Resolution Image Synthesis with Latent Diffusion Models," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 10684–10695.
- [9] Xiaosong Wang, Yifan Peng, Le Lu, Zhiyong Lu, Mohammadhadi Bagheri, and Ronald M Summers, "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*, 2017, pp. 2097–2106.
- [10] Takeru Miyato and Masanori Koyama, "cGANs with Projection Discriminator," in *International Conference* on Learning Representations, 2018.
- [11] Aditya Kuppa, Lamine Aouad, and Nhien-An Le-Khac, "Towards Improving Privacy of Synthetic DataSets," in *Annual Privacy Forum*. Springer, 2021, pp. 106–119.
- [12] Pranav Rajpurkar et al., "CheXNet: Radiologist-Level Pneumonia Detection on Chest X-Rays with Deep Learning," arXiv preprint arXiv:1711.05225, 2017.
- [13] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger, "Densely Connected Convolutional Networks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 4700–4708.
- [14] 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.
- [15] Patrick Esser, Robin Rombach, and Björn Ommer, "Taming Transformers for High-Resolution Image Synthesis," in *Proceedings of the IEEE/CVF Conference* on Computer Vision and Pattern Recognition, 2021, pp. 12873–12883.
- [16] Diederik P Kingma and Jimmy Ba, "Adam: A Method for Stochastic Optimization," arXiv preprint arXiv:1412.6980, 2014.