



HAL
open science

Generative medical image anonymization based on latent code projection and optimization

Huiyu Li, Nicholas Ayache, Hervé Delingette

► **To cite this version:**

Huiyu Li, Nicholas Ayache, Hervé Delingette. Generative medical image anonymization based on latent code projection and optimization. IEEE International Symposium on Biomedical Imaging (ISBI 2025), Apr 2025, Houston (Texas), United States. hal-04913904

HAL Id: hal-04913904

<https://inria.hal.science/hal-04913904v1>

Submitted on 27 Jan 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

GENERATIVE MEDICAL IMAGE ANONYMIZATION BASED ON LATENT CODE PROJECTION AND OPTIMIZATION

Huiyu Li, Nicholas Ayache, Hervé Delingette

Inria, Epione Team, Université Côte d’Azur, Sophia Antipolis, France

ABSTRACT

Medical image anonymization aims to protect patient privacy by removing identifying information, while preserving the data utility to solve downstream tasks. In this paper, we address the medical image anonymization problem with a two-stage solution: latent code projection and optimization. In the projection stage, we design a streamlined encoder to project input images into a latent space and propose a co-training scheme to enhance the projection process. In the optimization stage, we refine the latent code using two deep loss functions designed to address the trade-off between identity protection and data utility dedicated to medical images. Through a comprehensive set of qualitative and quantitative experiments, we showcase the effectiveness of our approach on the MIMIC-CXR chest X-ray dataset by generating anonymized synthetic images that can serve as training set for detecting lung pathologies. Source codes are available at <https://github.com/Huiyu-Li/GMIA>.

Index Terms— Medical image anonymization, Identity-utility trade-off, Latent code optimization.

1. INTRODUCTION

Medical images often contain sensitive information that can be linked to individual patients [1], posing risks of privacy breaches and identity attack. Therefore, robust anonymization techniques are essential to safeguard patient confidentiality while preserving the clinical utility of the images.

With the rise of generative models like Generative Adversarial Networks (GANs), new approaches have emerged that leverage the generative power of these networks to tackle anonymization by creating synthetic data. This approach ensures that while the synthetic data retains key statistical properties and utility for analysis, it no longer corresponds to any real individual, offering a robust solution for protecting sensitive information. For instance, Jeon et al. [2] introduced *k*-SALSA, a framework utilizing a GAN-based approach to synthesize a *K*-anonymous dataset from private retinal images. However, the process of sample aggregation in the latent space comes with the drawback of reducing the dataset size by a factor of *k*. To address this issue, Pennisi et al. [3] introduced a latent space navigation strategy to generate a diverse range of images. However, due to the weak supervision

in the navigation process, the clinical interpretability of the generated images remains limited.

Unlike existing generative medical image anonymization methods that typically rely on GAN inversion approach (“ \mathcal{E} -training”), where the encoder is trained with a fixed pre-trained StyleGAN generator to discover the latent representation of an input image, we approach the learning of the latent code as an image reconstruction task. This involves using a custom encoder and a co-training scheme to achieve a precise projection of the input image in its latent space. Furthermore, we adopt the latent code optimization technique from [4] to explicitly manage the trade-off between identity protection and data utility, incorporating deep identity and utility losses tailored specifically for medical images.

2. PROPOSED METHOD

Given a private real image dataset, we aim to generate an anonymized version that ensures patient privacy while preserving data utility. As shown in Fig. 1, we first project the real image into its latent space W through image reconstruction. Next, we initialize the anonymized latent code W_A as W and optimize it using two loss functions: an identity loss \mathcal{L}_{id} to obscure identifiable features, and a utility loss \mathcal{L}_{ut} to retain critical diagnostic information.

2.1. Latent Code Projection

We formulate the latent code projection as an image reconstruction task using an AE-GAN (Autoencoder-Generative Adversarial Network) architecture (see top in Fig. 1), which combines a custom encoder \mathcal{E} , a StyleGAN2 [5] generator \mathcal{G} and discriminator \mathcal{D} . The encoder transforms an real image X_R into its latent code $W = \mathcal{E}(X_R)$, while the generator reconstructs an approximation of the real image, $\hat{X}_R = \mathcal{G}(W)$, from the latent code. A discriminator $\mathcal{D}(X_R)$ is also employed to distinguish between real and generated images.

The custom encoder \mathcal{E} (lower left of Fig. 1) mirrors the generator \mathcal{G} , leveraging the inherent symmetry of the AE architecture. \mathcal{E} begins with an operation (fromIMG) that converts the input image into feature maps for subsequent convolutional processing. At each resolution, we implement a residual network architecture, consisting of two 3×3 convolution layers followed by a bilinear downsampling operation to

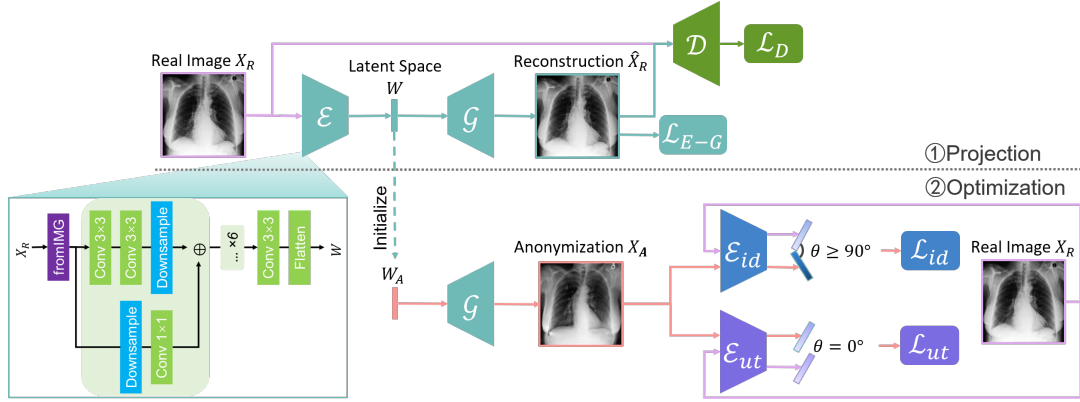


Fig. 1. Overview of the proposed method, consisting of two key stages: (1) AE-GAN network for latent code projection, and (2) Latent code optimization using identity removal loss and utility-preserving losses.

reduce the feature map size. The skip connection includes a bilinear downsampling operation and a 1×1 convolution layer. After 6 operation groups, \mathcal{E} projects the input image into the W latent space of StyleGAN2 [5].

To learn a more accurate mapping from the input image to its W latent space, we propose a co-training scheme to optimize the AE-GAN network. In each training step, the encoder \mathcal{E} and generator \mathcal{G} are jointly optimized using a composite loss function $\mathcal{L}_{\mathcal{E},\mathcal{G}}$, which integrates three key components: pixel-wise similarity loss, perceptual similarity loss, and adversarial (generator) loss. This is followed by optimizing the discriminator \mathcal{D} using the discriminator loss $\mathcal{L}_{\mathcal{D}}$.

2.2. Latent Code Optimization

Latent code optimization (see bottom right in Fig. 1) drives the anonymization process, with W_A as the only trainable parameter and initialized from W , while all other networks remain pre-trained and fixed. W_A is optimized using two deep loss functions: \mathcal{L}_{id} , which ensures that identifiable features are effectively obfuscated or removed, and \mathcal{L}_{ut} , which preserves critical visual and diagnostic information.

The identity loss \mathcal{L}_{id} ensures that X_A has a different identity from X_R , up to a desired margin. It is defined as follows:

$$\mathcal{L}_{id}(X_R, X_A) = \max(0, \cos(\mathcal{E}_{id}(X_R), \mathcal{E}_{id}(X_A)) - m) \quad (1)$$

where $\cos(a, b) = \frac{a \cdot b}{\max(\|a\|_2, \|b\|_2, \epsilon)}$ denotes the cosine similarity, and ϵ is a small value to avoid division by zero. \mathcal{E}_{id} is the identity encoder, which is pre-trained for identity recognition [6]. m is a hyperparameter that controls the dissimilarity between the real and the anonymized images, constrained to be equal to or greater than $\arccos(m)$.

The utility loss \mathcal{L}_{ut} enforces that utility attributes of X_R are preserved in X_A . It is defined as follows:

$$\mathcal{L}_{ut}(X_R, X_A) = \|\mathcal{E}_{ut}(X_R) - \mathcal{E}_{ut}(X_A)\|_2 \quad (2)$$

where \mathcal{E}_{ut} is the utility encoder, which is pre-trained to perform a downstream task (e.g. lung pathology classification).

3. EXPERIMENTS

Dataset. We perform anonymization on MIMIC-CXR-JPG dataset [7], using the default dataset partition. To ensure sufficient data per patient, we apply a filtering strategy (≥ 20 images per patient). This results in 50875/522/1641 images, corresponding to 1538/14/47 patients in the training/validation/testing sets, respectively. The utility task focuses on classifying lung pathology labels, while the identity classifier is designed to predict patient IDs.

Latent Code Projection. We evaluate the projection process using both the proposed "Co-training" scheme and the classic " \mathcal{E} -training" approach, comparing their performance in terms of image reconstruction fidelity.

Utility Preservation. To assess the preservation of utility attributes, we train utility classifiers on the training sets of real images X_R , reconstructed images \hat{X}_R , and anonymized images X_A , then evaluate them on the same unseen real test set.

Identity Anonymization. To measure the privacy-preserving properties of our approach, we first assess the singling-out risk¹ on testing set with two metrics: HR (hidden rate) and LC (local cloaking). HR refers to the percentage of individuals in the real dataset whose anonymized images is not the most similar to them using the identity embedding. LC is defined as the median number of images ranked between a real image and its anonymization according to the identity distance.

Then, we introduce a protocol to assess linkability risk² from two perspectives: **Inner risk:** Verifying whether two anonymized images belong to the same patient. **Outer risk:** Verifying whether a real and its anonymized image belong to the same patient. Specifically, we frame the evaluation as a standard identity verification task, where the identity classifier is trained on the real dataset to determine whether two images belong to the same patient. For the inner risk, the test set

¹Singling out: the possibility to isolate some or all records which identify an individual in the dataset.

²Linkability: the ability to link, at least, two records concerning the same data subject or a group of data subjects.

consists of image pairs from the anonymized test set. For the outer risk, the test set comprises image pairs, each containing a real image and its anonymized version from the test set.

Table 1. The results of image reconstruction. 'PSNR': peak signal-to-noise ratio, 'SSIM': structural similarity, 'IW_SSIM': information content weighted structural similarity, 'LPIPS': learned perceptual image patch similarity.

Method	PSNR \uparrow	SSIM \uparrow	IW_SSIM \uparrow	LPIPS \downarrow
Co-training	24.319	0.968	0.602	0.340
\mathcal{E} -training	19.477	0.936	0.363	0.446

Lastly, we perform membership inference attacks (MIA) [8] to predict whether an anonymized image was part of the real training set (given the model outputs of the anonymized image). The attacker is trained on the real dataset and evaluated on the anonymized dataset (target dataset).

4. RESULTS

Image Reconstruction. Table 1 shows that the co-training scheme outperforms the \mathcal{E} -training approach in reconstruction accuracy, as indicated by higher PSNR and IW_SSIM scores, highlighting the benefits of joint optimization for enhancing reconstruction quality. The results are further validated by the visual comparison in Fig. 2, where the co-training scheme recovers input images with finer details and higher fidelity. In contrast, the \mathcal{E} -training scheme exhibits noticeable discrepancies when compared to the original images.




Fig. 2. Reconstruction results. The first row displays the real images X_R . The last two rows show the reconstructed images \hat{X}_R produced by the proposed co-training scheme and the \mathcal{E} -training scheme, respectively.

Utility Preservation. As shown in Table 2, the performance of classifiers trained on anonymized datasets (row 4-6) exhibits minimal difference compared to the one trained on real dataset (row 1), indicating that the anonymized datasets maintain a similar level of data utility compared to their real counterparts. Moreover, incorporating \mathcal{L}_{ut} (row 4, 6) yields better results compared to those without it (row 2, 5), highlighting the effectiveness of the utility loss \mathcal{L}_{ut} in preserving utility-

related information during optimization. Additionally, classifier trained on reconstructed datasets using \mathcal{E} -training scheme (row 3) shows the weakest performance, highlighting the advantages of the proposed co-training scheme.

Table 2. Results of utility preservation. 'Acc': accuracy, 'AP': average precision, 'AUROC': area under the ROC curve, 'F1': F1 score.

Training Data	Acc \uparrow	AP \uparrow	AUROC \uparrow	F1 \uparrow	
X_R	0.740	0.631	0.774	0.511	
\hat{X}_R	Co-training	0.718	0.586	0.737	0.451
	\mathcal{E} -training	0.654	0.472	0.640	0.307
X_A	\mathcal{L}_{ut}	0.726	0.600	0.747	0.493
	\mathcal{L}_{id}	0.696	0.559	0.713	0.424
	$\mathcal{L}_{ut} + \mathcal{L}_{id}$	0.713	0.578	0.737	0.531

Identity Anonymization. Table 3 denotes that the anonymized images (row 3-5, 7-9) exhibit a lower linkability risk compared to their real counterparts (row 1). Additionally, the anonymized datasets exhibit high inner risk and a low outer risk, which is expected. The high inner risk indicates that the anonymized images retain relative likability relationships (e.g., images from the same patient remain associated with each other), while the low outer risk signifies that the linkability relationship between anonymized images and their real counterparts is effectively removed (e.g., anonymized images do not resemble their real counterparts in terms of identity).

Table 3. Results of linkability risk. 'Pre': pre-training, 'F1': F1 score, 'Acc': accuracy, 'TAR': true accept rate, 'FAR': false accept rate.

	Testing Data	F1 \uparrow	Acc \uparrow	TAR \uparrow	FAR \downarrow	
	X_R	0.983	0.983	0.966	0.034	
Inner risk	\hat{X}_R Co-training	0.611	0.720	0.441	0.559	
		\mathcal{L}_{ut}	0.665	0.748	0.500	0.500
	X_A	\mathcal{L}_{id}	0.855	0.873	0.746	0.254
		$\mathcal{L}_{ut} + \mathcal{L}_{id}$	0.887	0.899	0.798	0.202
Outer risk	\hat{X}_R Co-training	0.139	0.075	0.075	0.925	
		\mathcal{L}_{ut}	0.136	0.073	0.073	0.927
	X_A	\mathcal{L}_{id}	0.000	0.000	0.000	1.000
		$\mathcal{L}_{ut} + \mathcal{L}_{id}$	0.000	0.000	0.000	1.000

Moreover, incorporating \mathcal{L}_{id} increases the inner risk while reducing the outer risk. This suggests the loss \mathcal{L}_{id} helps push the anonymized images belonging to the same patient cluster together while breaking the linkability to their real counterparts. Similar results are shown in Table 4, where \mathcal{L}_{id} plays a crucial role in reducing the singling out risk, as incorporating \mathcal{L}_{id} (row 3-4) produces anonymized datasets with a low risk of singling out compared to the datasets without it (row 1-2).

In the MIA attack, attackers achieve high accuracy (Acc: 0.968) during training, demonstrating their effectiveness on the test dataset. However, during the evaluation stage, their

Table 4. Results for singling out risk on the 1641 images. ' N_R ': median number of real images, ' N_A ' median number of anonymized images.

Testing Data		HR \uparrow	LC \uparrow	
			N_R	N_A
\hat{X}_R	Co-training	0.936	41	58
	\mathcal{L}_{ut}	0.936	41	56
X_A	\mathcal{L}_{id}	1.000	1546	1639
	$\mathcal{L}_{ut} + \mathcal{L}_{id}$	1.000	1529	1639

performance significantly drops, with an accuracy of 0. This indicates a considerable challenge in correctly predicting the membership status of anonymized images.




Fig. 3. Anonymization results. Real images X_R randomly selected from the training, validation, and test sets are displayed in the first column. The corresponding reconstructed images \hat{X}_R are displayed in the second column. The anonymized images X_A are displayed in the last three columns.

Qualitative Results The visualization results are shown in Fig. 3. In this figure, the anonymized images optimized using only the utility loss \mathcal{L}_{ut} (column 3) exhibit greater visual similarity to their real counterparts (column 1), while those optimized using only the identity loss \mathcal{L}_{id} (column 4) appear more distinct from their originals. Additionally, the anonymized images optimized with both the \mathcal{L}_{ut} and the \mathcal{L}_{id} (column 5) strike a balance, appearing more realistic by simultaneously considering both identity removal and utility preservation.

5. CONCLUSIONS

In this paper, we introduce an innovative two-stage approach to address medical image anonymization. First, we precisely project the input image into its latent space using image reconstruction with a custom encoder and an effective co-training scheme. Then, we optimize the latent code to remove identifiable information while preserving utility using two deep loss functions. Our results demonstrate that the method effectively anonymizes the images while better preserving

utility attributes, leading to improved de-identification and retention of critical clinical information.

6. REFERENCES

- [1] 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 x-ray data,” *Scientific Reports*, 2022.
- [2] Minkyu Jeon, Hyeonjin Park, Hyunwoo J Kim, Michael Morley, and Hyunghoon Cho, “k-salsa: k-anonymous synthetic averaging of retinal images via local style alignment,” in *European Conference on Computer Vision*. Springer, 2022.
- [3] Matteo Pennisi, Federica Proietto Salanitri, Giovanni Bellitto, Simone Palazzo, Ulas Bagci, and Concetto Spampinato, “A privacy-preserving walk in the latent space of generative models for medical applications,” in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2023.
- [4] Simone Barattin, Christos Tzelepis, Ioannis Patras, and Nicu Sebe, “Attribute-preserving face dataset anonymization via latent code optimization,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023.
- [5] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila, “Analyzing and improving the image quality of stylegan,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020.
- [6] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou, “Arcface: Additive angular margin loss for deep face recognition,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019.
- [7] Alistair Johnson, Matt Lungren, Yifan Peng, Zhiyong Lu, Roger Mark, Seth Berkowitz, and Steven Horng, “Mimic-cxr-jpg-chest radiographs with structured labels,” *PhysioNet*, 2019.
- [8] Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov, “Membership inference attacks against machine learning models,” in *2017 IEEE symposium on security and privacy (SP)*. IEEE, 2017.

Acknowledgements. This work has been supported by the French government, through the 3IA Côte d’Azur Investments in the Future project managed by the National Research Agency (ANR) with the reference number ANR-19-P3IA-0002, and supported by the Inria Sophia Antipolis - Méditerranée, “NEF” computation cluster.