

MedSynth: Leveraging Generative Model for Healthcare Data Sharing

Renuga Kanagavelu¹[0000-0002-2808-3093], Madhav Walia²[0009-0001-3943-8136],
Yuan Wang¹[0000-0002-5258-1500], Huazhu Fu¹[0000-0002-9702-5524], Qingsong
Wei¹[0000-0003-4570-7328], Yong Liu¹[0000-0002-1590-2029], and Rick Siow Mong
Goh¹[0000-0001-9116-1595]

¹ Institute of High Performance Computing (IHPC), Agency for Science, Technology
and Research (A*STAR), Singapore

² SRM University, AP, India

{renuga_k, wang_yuan, wei_qingsong, liuyong, gohsm}@ihpc.a-star.edu.sg,
hzfu@ieee.org, madhav_walia@srmmap.edu.in

Abstract. Sharing medical datasets among healthcare organizations is essential for advancing AI-assisted disease diagnostics and enhancing patient care. Employing techniques like data de-identification and data synthesis in medical data sharing, however, comes with inherent drawbacks that may lead to privacy leakage. Therefore, there is a pressing need for mechanisms that can effectively conceal sensitive information, ensuring a secure environment for data sharing. Dataset Condensation (DC) emerges as a solution, creating a reduced-scale synthetic dataset from a larger original dataset while maintaining comparable training outcomes. This approach offers advantages in terms of privacy and communication efficiency in the context of medical data sharing. Despite these benefits, traditional condensation methods encounter challenges, particularly with high-resolution medical datasets. To address these challenges, we present *MedSynth*, a novel dataset condensation scheme designed to efficiently condense the knowledge within extensive medical datasets into a generative model. This facilitates the sharing of the generative model across hospitals without the need to disclose raw data. By combining an attention-based generator with a vision transformer (ViT), *MedSynth* creates a generative model capable of producing a concise set of representative synthetic medical images, encapsulating the features of the original dataset. This generative model can then be shared with hospitals to optimize various downstream model training tasks. Extensive experimental results across medical datasets demonstrate that *MedSynth* outperforms state-of-the-art methods. Moreover, *MedSynth* successfully defends against state-of-the-art Membership Inference Attacks (MIA), highlighting its significant potential in preserving the privacy of medical data.

Keywords: Data sharing · Privacy · Dataset Condensation · Generative Model · Membership Inference Attack · Vision Transformer.

1 Introduction

Incorporating Artificial Intelligence (AI) into medical practice leads to advancement and breakthroughs in areas like radiology, pathology, and the overall healthcare ecosystem [1, 2]. However, the data-intensive nature of AI in medicine emphasizes the need for substantial and high-quality data, highlighting the significance of collaborative medical data sharing among hospitals [3]. Nevertheless, efficient data sharing faces challenges due to privacy concerns. Using methods such as data synthesis [4] and data de-identification [5] in medical data sharing, however, poses inherent challenges that could potentially result in privacy leakage. Data de-identification techniques always carry a risk of re-identification [6], and there is evidence that the models used in data synthesis may leak information related to the training samples, making them prone to Membership Inference Attacks (MIA) [7]. Hence, it is essential to implement secure and dependable solutions for sharing healthcare data.

Dataset Condensation (DC) [8] involves creating a reduced-scale synthetic dataset derived from a large original dataset while still achieving comparable training outcomes. This method provides advantages in terms of privacy [9, 10] and data storage efficiency [11]. Therefore, it is beneficial to explore medical dataset condensation to address current challenges associated with medical dataset sharing. Nevertheless, when dealing with high-resolution medical datasets, traditional condensation methods [12, 13, 14, 15, 11] encounter difficulties as they directly extract information from the original dataset into pixel space, and the feature distribution of condensed samples frequently lacks diversity.

To address the challenges, we present a novel dataset condensation (DC) approach named *MedSynth* aiming to condense the knowledge of a large medical dataset into a single generative model. With this generative model, one can generate a concise set of representative synthetic images, encapsulating the features of the original dataset. This is achieved by integrating an attention-based generator with vision transformer (ViT)-based feature refinement. By sharing only the generative model, a hospital or other medical facility can securely contribute its local data knowledge to another hospital. This facilitates enhancing data resources for downstream tasks such as training and fine-tuning, while effectively mitigating potential privacy risks.

We evaluate the effectiveness of *MedSynth* using ISIC 2019 and Alzheimer’s MRI datasets and compare its efficiency with state-of-the-art methods. The key contributions of our work can be outlined as follows:

1. We propose *MedSynth*, a novel dataset condensation scheme that efficiently condenses the knowledge within extensive medical datasets into a generative model for secure sharing across hospitals without disclosing raw data and ensuring privacy protection.
2. We enhance the extraction of fine-grained information from medical images by combining an attention-based generator with a ViT for feature matching.
3. Extensive experimentation on medical datasets demonstrates that *MedSynth* outperforms state-of-the-art methods.

4. We conduct a membership inference attack on *MedSynth*'s generative model to confirm its resilience and successful protection of medical data privacy.

2 Related Work

Deep learning models need to be trained on extensive datasets for accurate disease diagnosis, emphasizing the significance of healthcare data sharing. De-identification [5] is a data-sharing technique that focuses on removing or making anonymous any personally identifiable information from medical records. However, there is consistently a potential risk of re-identification in this process [6]. An alternative solution to address data sharing limitations involves using Generative Adversarial Networks (GANs) [4]. It facilitates the generation of anonymous and potentially limitless synthetic datasets. Nonetheless, there is evidence [7] that the models used in data synthesis may leak information related to the training samples. Federated learning offers an alternative that trains a joint model across many hospitals without sharing raw data [16, 17]. Instead, each hospital trains the model locally and shares only its updates, keeping data private. Nonetheless, it introduces a vulnerability, as model updates sent to the central server are susceptible to gradient leakage attacks [18], potentially allowing the reconstruction of original data.

A recent strategy called dataset condensation has emerged to address these limitations. It can be achieved in two ways: either condensing the knowledge of an entire original dataset into a few synthetic images or into a generative model. Several techniques follow the first method [12, 13, 14, 15, 11]. These methods begin with a small number of learnable image tensors, which are then updated by comparing the training trajectories [15], embedding distributions [13, 14], or training gradients [11] with the original images. The second approach is utilized by [19, 20]. IT-GAN [19] aims to explore if a fixed GAN can generate informative images without changing the dataset size or reducing training costs. DiM [20], on the other hand, condenses the information of the original dataset into a generator, although it has limitations when dealing with high-resolution medical images. Current medical dataset distillation methods [21],[22] consider learning soft labels in the synthetic set, meaning that the label is trainable for improved information compression. The researchers [23] also proposed a distillation-based approach, aiming to match the parameters of the teacher networks trained on the original dataset with the student networks built on the distilled dataset. In [24], a generalizable dataset distillation-based federated learning (GDD-FL) framework is proposed to achieve communication-efficient federated skin lesion classification. Our method is a generative-based method that condenses the knowledge within extensive medical datasets into a generative model.

3 Methodology

The *MedSynth* works in two phases as shown in Figure 1. In the first phase, an attention-based generator is trained to capture the features of the original

medical dataset. In the second phase, the generator is fine-tuned using a ViT to condense the important features into the synthetic dataset. Fine-tuning is accomplished by employing logit matching to compare the features of the synthetic dataset with those of the original dataset.




Fig. 1. Workflow of our framework: 1) Phase I- Pre-train the attention-based generator to create meaningful correspondences with original images. 2) Phase II- Fine-tuning the generator by applying logit matching on image embeddings of original and synthetic datasets produced by a fine-tuned ViT.

3.1 Phase I - Pre-Train the Generator

The attention-based generator used in the framework comprises an encoder, a transition layer, and a decoder. The encoder reduces the image’s feature maps to a quarter of their original size. Subsequently, in the transition layer, we replaced the inner convolutional layers with six residual attention blocks (RABs). This modification effectively captures long-range dependencies and highlights latent features within the medical images. Finally, the decoder performs up-sampling to generate two outputs: a generated image, denoted as g , and an attention mask, denoted as a . The synthetic image, denoted as s , is defined as the weighted combination of the generated image g and input x by a .

$$s = a \times x + (1 - a) \times g \quad (1)$$

The synthetic image s is then fed to the discriminator F_a , which comprises of 9 convolution layers and 2 fully connected layers. As described in [25], Wasserstein distance with gradient penalty is utilized for calculating the GAN loss, defined as:

$$L_G = E_x[F_a(s)] - E_x[F_a(x)] + \lambda_{GP} (\nabla_z F_a(z)^2 - 1)^2 \quad (2)$$

$$z = \epsilon \cdot x + (\epsilon - 1) \cdot s, \quad \epsilon \sim \mathcal{U}[0, 1] \quad (3)$$

where $E_x[\cdot]$ is the expected value of a function the input data and ϵ is the same size as the input x , and has random values from 0 to 1. The Wasserstein

distance is evaluated by the first two terms in L_G , while the last term accounts for the gradient penalty [26]. Additionally, we replaced standard batch normalization with spectral normalization in the generator, as recommended by [27], for improved stability during training with large batches and highly textured data like medical images. The training phase is crucial, as our experimental findings indicate a significant 25% degradation in classification performance without this training step.

3.2 Phase II - Fine-Tune the Generator using ViT

In the second phase, the pre-trained generator from the first phase is fine-tuned using a ViT to extract the important features into the synthetic dataset. While previous works [13, 11, 20, 28], have utilized ConvNet models for feature vector extraction, our study opts for the ViT architecture [29]. Unlike ConvNet, ViTs utilize self-attention mechanisms, allowing them to directly compare and relate features across the entire image and capture the fine-grained details in medical scans. Their inherent inductive bias towards global features enables them to learn efficiently from limited data, potentially overcoming the data bottlenecks that hinder CNN-based approaches in this domain.

ViT for feature extraction - In our approach, we initially fine-tuned the ViT on the original dataset to facilitate the extraction of feature vectors. Subsequently, the feature vectors of both the original and synthetic images, extracted from the ViT, are compared using logit matching. This process aims to condense the essential attributes of the original images into the generator, significantly enhancing the quality of synthetic images produced and serving as a key process in the dataset condensation.

However, due to data scarcity and class imbalance in medical datasets, fine-tuning the ViT becomes a challenge as the model starts to overfit. Our work mitigates this problem by applying Low-Rank Adaptation (LoRA) [30] to perform parameter-efficient fine-tuning. LoRA significantly diminishes the number of trainable parameters, reducing the risk of overfitting when fine-tuning a large transformer-based model with limited data points. This enables the utilization of ViT with medical datasets, which are typically limited in size.

Fine-tune generator via logit matching - For a batch consisting of authentic images x and generated images S , we utilize a network n to predict the classification logits for these images. To formalize, the alignment of logits can be articulated as follows:

$$L_n = MSE(n(S); n(x)) \quad (4)$$

Here MSE is the Mean Squared Error loss over the two feature vectors produced by the model. The Logit Matching loss L_n seeks to minimize the prediction logits, which directly impact the outcomes of subsequent classification tasks. Finally, the total can be written as:

$$L_{total} = L_n + \lambda \cdot L_g \quad (5)$$

This generative model can be shared securely, allowing hospitals to enhance their resources without compromising privacy.

4 Experimental Analysis

4.1 Datasets

Our evaluations were performed using the Alzheimer’s Disease (AD) [31] and the International Skin Imaging Collaboration (ISIC) 2019 dataset [32]. The Alzheimer’s Disease dataset is a collection of MRI images. The images are classified into four distinct categories: Mild Demented, Moderate Demented, Non-Demented, and Very Mild Demented. It has 5121 train images and 1279 test images.

The ISIC 2019 dataset comprises 20,264 train images and 5,067 test images available for the classification of dermoscopic images among nine different diagnostic categories: Melanoma, Melanocytic nevus, Basal cell carcinoma, Actinic keratosis Benign keratosis (solar lentigo / seborrheic keratosis/lichen planus-like keratosis), Dermatofibroma, Vascular lesion, Squamous cell carcinoma and none of the above.

4.2 Implementation Details

Experimental Setup The training for the generator occurs in two phases (refer to Fig. 1). In the first phase, we train the attention-based generator. The default epoch size is set to 150. The momentum for the ADAM optimizer is set to 0.9. The learning rate for *MedSynth* generator is set to 0.0001 while for the ViT in the second phase is set to 0.01. The batch size B is kept at 32 to avoid memory overflow. The implementation utilizes the PyTorch framework and is run on 4 Nvidia V100 GPUs.

Performance Assessment and Metrics To assess the effectiveness of the *MedSynth*, we select a pre-trained ViT as the backbone for the downstream classification task. We first fine-tune this ViT using LoRa on the condensed datasets obtained with MedSynth and the baseline methods, utilizing a learning rate of 0.01, and a batch size of 32. Subsequently, we evaluate and report the resulting performance in terms of the area under the Receiver Operating Characteristic curve (AUC) and accuracy.

4.3 Comparisons with State-Of-The-Art

In Table 1, we present a comparison of our method results with other state-of-the-art distribution-based and GAN-based condensation methods. We compare our results with DiM [20] that uses vanilla conditional GAN. For the other methods, a ResNet18 architecture is used, whereas our work utilizes ViT-base16 with LoRA as described in [29, 30]. Both the architectures are pre-trained on ImageNet. The

experimental findings across the ISIC 2019 and Alzheimer datasets demonstrate that *MedSynth* outperforms baseline methods, reducing the dataset size by 95% while maintaining a close approximation of 96% of the original classification performance in terms of AUC on ISIC 2019 and Alzheimer datasets.

We believe two major factors contribute to this result: 1) The presence of residual attention blocks in the generator, aiding in capturing long-range dependencies, thus facilitating effective dataset condensation into the generator. 2) ViT’s self-attention mechanism enables the model to capture global interactions among all components of the input image, coupled with stronger generalization ability through pre-trained weights, further improves the quality of the generated condensed datasets.

We compare our method with non-condensation methods, specifically using deep convolutional GAN (DCGAN) [33]. We achieved a performance improvement of about 7% on the Alzheimer dataset and about 3% on the ISIC 2019 dataset compared with DCGAN.

We consider the condensation ratio which is the ratio of the condensed dataset size to the whole original training dataset size as a metric to measure the size of the synthetic dataset required to achieve a similar performance on the whole original dataset. In comparison to training on the entire dataset, using only 50 condensed images per class of Alzheimer’s and ISIC 2019 synthetic data resulted in approximately 96% of the original classification performance. Furthermore, a dataset comprising 25,331 skin lesion (ISIC) images typically demands around 10 GB of storage. In contrast, the generative model derived from it only requires about 524 MB, reducing the communication cost as well as the storage cost by a factor of 20. This illustrates how a generative model sharing can improve efficiency in healthcare settings, enhancing data portability, sharing, and distribution without excessively burdening the backbone network and ensuring privacy protection.

4.4 Generalization Ability Comparison



We also assess the generalization capability of our approach across various deep networks, including ConvNet, ResNet18, and DenseNet, and compare the performance with the DiM method. The findings are shown in Fig. 2. AUC scores for various models trained on the condensed Alzheimer’s dataset with images per class = 50 are compared. Our method demonstrates outstanding generalization performance across various architectures. This highlights that the condensed images produced by *MedSynth* can be utilized to train different networks, making it easier for users to select a model as required by their application.

4.5 Membership Inference Attack Analysis

We evaluate MedSynth’s resiliency against the MIA attack. We generate the white-box and black-box attacks based on the threat model proposed in [34]. For both, the attacker training set consists of a random 10% of the original dataset (ISIC/Alzheimer) with synthetic fake samples as non-members. In the white box

Table 1. Performance Evaluation with varying images per class (IPC) on ISIC 2019 and Alzheimer datasets.

Method	IPC	Dataset					
		Alzheimer's			ISIC 2019		
		Cond Ratio%	AUC Score	Accuracy	Cond Ratio%	AUC Score	Accuracy
DM [14]	50	0.98	65.90	79.34	0.27	77.90	41.24
	100	1.95	71.73	74.07	0.49	80.65	44.46
IT-GAN	50	0.98	78.84	79.54	0.27	77.24	77.62
	100	1.95	82.11	83.48	0.49	80.53	81.03
DiM	50	0.98	89.35	89.46	0.27	81.87	82.33
	100	1.95	90.28	91.13	0.49	82.94	83.25
MedSynth	50	0.98	94.90	95.58	0.27	83.12	85.54
	100	1.95	96.27	97.11	0.49	84.93	87.22
DCGAN		-	90.07	91.18	-	82.57	84.92
Original Dataset		-	97.19	98.83	-	86.19	88.95

**Fig. 2.** Generalization ability comparison to DiM.**Fig. 3.** Membership Inference Attack -Accuracy on the Generator

attack, knowing the target GAN architecture, the attacker inputs the training set to target GAN's discriminator, extracts and sorts the prediction probabilities and uses the highest probabilities to predict the training set members. In the black-box attack, without knowing the target GAN architecture, the attacker first trains a local GAN using the target GAN samples and carry out the steps typical of a white-box attack. The attack's accuracy is determined by the percentage of correctly identified images from the training set. Lower attack accuracy, below random guessing, indicates increased model security against MIA attacks. We compare the accuracy of both scenarios to a baseline corresponding to random guesses made by a third party on the membership of samples in the dataset. As shown in Fig. 3, we see the accuracy for both scenarios lies under the baseline, proving the generator is safe against MIA. This happens because the generator is distilled with knowledge about a condensed version of the dataset, which makes it difficult to reverse-engineer individual samples from the original dataset.

5 Conclusion

In this work, we proposed a novel generative model-based condensation technique to improve the high-resolution medical dataset’s condensation process using an attention-based generator combined with ViT-based feature refinement. Incorporating DC with specialized GAN architectures for medical images and attention-based foundation models is more effective in condensing medical images. We evaluated the effectiveness of our proposed approach on several healthcare datasets and achieved an AUC score similar to the original dataset. We also looked at how these generative models are safe from membership inference attacks, making them safe to use in the medical domain. By sharing the generative models instead of the entire raw images, it lowers the amount of storage and bandwidth needed for data storage and transfer. Our evaluation shows that the condensed medical datasets/generative models obtained with our method could be more securely and efficiently shared among healthcare facilities.

Acknowledgments. This Research is supported by the RIE2025 Industry Alignment Fund – Industry Collaboration Project (IAF-ICP) (Award No: I2301E0020), administered by A*STAR. This work is supported by A*STAR Central Research Fund “A Secure and Privacy Preserving AI Platform for Digital Health”.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

- [1] Y. Chen et al. “Generative Adversarial Networks in Medical Image augmentation: A review”. In: *Comput Biol Med* 144 (2022), p. 105382.
- [2] Jieneng Chen et al. *3D TransUNet: Advancing Medical Image Segmentation through Vision Transformers*. 2023. arXiv: 2310.07781 [cs.CV].
- [3] Mana Azarm-Daigle, Craig Kuziemy, and Liam Peyton. “A Review of Cross Organizational Healthcare Data Sharing”. In: *Procedia Computer Science* 63 (Dec. 2015), pp. 425–432.
- [4] Aldren Gonzales, Guruprabha Guruswamy, and Scott Smith. “Synthetic data in health care: A narrative review”. In: *PLOS Digital Health* 2 (Jan. 2023), e0000082.
- [5] Khaled El Emam, Sam Rodgers, and Bradley Malin. “Anonymising and sharing individual patient data”. In: *BMJ* 350 (2015).
- [6] Kerstin N. Vokinger, Daniel J. Stekhoven, and Michael Krauthammer. “Lost in Anonymization — A Data Anonymization Reference Classification Merging Legal and Technical Considerations”. In: *The Journal of Law, Medicine & Ethics* 48.1 (2020), pp. 228–231.
- [7] Z. Zhang, C. Yan, and B. A. Malin. “Membership inference attacks against synthetic health data”. In: *J Biomed Inform* 125 (2022), p. 103977.
- [8] Tongzhou Wang et al. “Dataset Distillation”. In: *CoRR* abs/1811.10959 (2018). arXiv: 1811.10959. URL: <http://arxiv.org/abs/1811.10959>.

- [9] Tian Dong, Bo Zhao, and Lingjuan Lyu. *Privacy for Free: How does Dataset Condensation Help Privacy?* 2022. arXiv: 2206.00240 [cs.CR].
- [10] Dingfan Chen, Raouf Kerkouche, and Mario Fritz. *Private Set Generation with Discriminative Information*. 2022. arXiv: 2211.04446 [cs.CR].
- [11] Bo Zhao, Konda Reddy Mopuri, and Hakan Bilen. *Dataset Condensation with Gradient Matching*. 2021. arXiv: 2006.05929 [cs.CV].
- [12] Bo Zhao and Hakan Bilen. “Dataset condensation with differentiable siamese augmentation”. In: *International Conference on Machine Learning*. PMLR. 2021, pp. 12674–12685.
- [13] Kai Wang et al. *CAFE: Learning to Condense Dataset by Aligning Features*. 2022. arXiv: 2203.01531 [cs.CV].
- [14] Bo Zhao and Hakan Bilen. *Dataset Condensation with Distribution Matching*. 2022. arXiv: 2110.04181 [cs.LG].
- [15] George Cazenavette et al. *Dataset Distillation by Matching Training Trajectories*. 2022. arXiv: 2203.11932 [cs.CV].
- [16] H. Brendan McMahan et al. *Communication-Efficient Learning of Deep Networks from Decentralized Data*. 2023. arXiv: 1602.05629 [cs.LG].
- [17] Yuan Wang et al. “An Aggregation-Free Federated Learning for Tackling Data Heterogeneity”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2024, pp. 26233–26242.
- [18] Yangsibo Huang et al. *Evaluating Gradient Inversion Attacks and Defenses in Federated Learning*. 2021. arXiv: 2112.00059 [cs.CR].
- [19] Bo Zhao and Hakan Bilen. *Synthesizing Informative Training Samples with GAN*. 2022. arXiv: 2204.07513 [cs.LG].
- [20] Kai Wang et al. *DiM: Distilling Dataset into Generative Model*. 2023. arXiv: 2303.04707 [cs.CV].
- [21] Guang Li et al. “Soft-Label Anonymous Gastric X-Ray Image Distillation”. In: *2020 IEEE International Conference on Image Processing (ICIP)* (2020), pp. 305–309.
- [22] Guang Li et al. “Compressed gastric image generation based on soft-label dataset distillation for medical data sharing”. In: *Computer Methods and Programs in Biomedicine* 227 (2022), p. 107189. ISSN: 0169-2607.
- [23] Guang Li et al. “Dataset Distillation for Medical Dataset Sharing”. In: *AAAI-23 Workshop on Representation Learning for Responsible Human-Centric AI* (2023).
- [24] Yuchen Tian et al. “Communication-Efficient Federated Skin Lesion Classification with Generalizable Dataset Distillation”. In: *MICCAI 2023 Workshops*. Vancouver, BC, Canada: Springer-Verlag, 2023. ISBN: 978-3-031-47400-2.
- [25] Euijin Jung, Miguel Luna, and Sang Hyun Park. “Conditional generative adversarial network for predicting 3d medical images affected by alzheimer’s diseases”. In: *Predictive Intelligence in Medicine: Third International Workshop, PRIME 2020, Held in Conjunction with MICCAI 2020*. Springer. 2020, pp. 79–90.

- [26] Ishaan Gulrajani et al. *Improved Training of Wasserstein GANs*. 2017. arXiv: 1704.00028 [cs.LG].
- [27] Takeru Miyato et al. *Spectral Normalization for Generative Adversarial Networks*. 2018. arXiv: 1802.05957 [cs.LG].
- [28] Bo Zhao and Hakan Bilen. *Dataset Condensation with Differentiable Siamese Augmentation*. 2021. arXiv: 2102.08259 [cs.LG].
- [29] Alexey Dosovitskiy et al. *An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale*. 2021. arXiv: 2010.11929 [cs.CV].
- [30] Edward J. Hu et al. *LoRA: Low-Rank Adaptation of Large Language Models*. 2021. arXiv: 2106.09685 [cs.CL].
- [31] Sarvesh Dubey. *Alzheimer’s Dataset (4 class of Images)*. 2020. URL: <https://www.kaggle.com/datasets/tourist55/alzheimers-dataset-4-class-of-images>.
- [32] *Skin Lesion Images for Melanoma Classification*. <https://www.kaggle.com/datasets/andrewmvd/isic-2019>.
- [33] Alec Radford, Luke Metz, and Soumith Chintala. *Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks*. 2016. arXiv: 1511.06434 [cs.LG].
- [34] Jamie Hayes et al. “LOGAN: Membership Inference Attacks Against Generative Models”. In: *Proceedings on Privacy Enhancing Technologies* 2019 (2017), pp. 133–152.