arXiv:2403.10573v2 [eess.IV] 7 Jul 2024

Medical Unlearnable Examples: Securing Medical Data from Unauthorized Training via Sparsity-Aware Local Masking

Weixiang Sun [1] **Yixin Liu** [2] **Zhiling Yan** [2] **Kaidi Xu** [3] **Lichao Sun** [2]

Abstract

The rapid expansion of AI in healthcare has led to a surge in medical data generation and storage, boosting medical AI development. However, fears of unauthorized use, like training commercial AI models, hinder researchers from sharing their valuable datasets. To encourage data sharing, one promising solution is to introduce imperceptible noise into the data. This method aims to safeguard the data against unauthorized training by inducing degradation in the generalization ability of the trained model. However, they are not effective and efficient when applied to medical data, mainly due to the ignorance of the sparse nature of medical images. To address this problem, we propose the Sparsity-Aware Local Masking (SALM) method, a novel approach that selectively perturbs significant pixel regions rather than the entire image as previously. This simple yet effective approach, by focusing on local areas, significantly narrows down the search space for disturbances and fully leverages the characteristics of sparsity. Our extensive experiments across various datasets and model architectures demonstrate that SALM effectively prevents unauthorized training of different models and outperforms previous SoTA data protection methods.

1. Introduction

The expansion of artificial intelligence (AI) in healthcare has significantly increased the production and storage of sensitive medical data (Rajkomar et al., 2018; Rasmy et al., 2021). This data plays a crucial role in advancing research in related fields (Zhang et al., 2022; Kelly et al., 2019; Rasmy et al., 2021). The more high-quality, open-source datasets that are available, the more contributions can be made by talented researchers to the development of the field (Johnson et al., 2019; Irvin et al., 2019; Pedrosa et al., 2019; Bilic et al., 2023; Rasmy et al., 2021). However, many dataset creators are reluctant to open-source their work due to concerns over unauthorized use, such as training commercial models (Edwards, 2022; Hill & Krolik, 2019).

Conventional image classification datasets are relatively easy to obtain and label by laypersons, as shown by ImageNet's use of online crowdsourcing (Deng et al., 2009; Zhang et al., 2021). But in the medical field, annotating is a much more complex process requiring specialized knowledge, usually undertaken by experts like radiologists (Kermany et al., 2018; Simpson et al., 2019; Kather et al., 2019). In addition, it is also necessary to describe the severity of the disease and its relationship with the surrounding tissues. This means that constructing a medical dataset is a labor-intensive task. Unauthorized use could lead to infringement of the creator's rights, leading to a reluctance to release further data publicly. Moreover, from a patient's perspective, concerns about their information being exploited for commercial purposes might decrease their willingness to authorize their data for research (Koh et al., 2011; Trinidad et al., 2020). Therefore, the construction and release of a medical dataset are not only time-consuming and labor-intensive but also fraught with significant ethical and privacy challenges. The reduction in quality open source datasets, resulting from both of the scenarios mentioned, in turn, slows down the development of medical AI (Alberdi et al., 2016; Forghani et al., 2019; Gunraj et al., 2020).

To defend unauthorized use and encourage sharing data, Huang et al. (2021) proposed the technique of contaminating data with imperceptible noise. Models trained on this noise-contaminated data exhibit poor performance for normal utilization. Specifically, these methods create "unlearnable" examples from clean data by adding imperceptible noise, demonstrating significant protective capabilities. The design of this error-minimizing noise is based on the intuitive idea that an example with a higher training loss could contain more information to learn. Consequently, this noise protects the data by minimizing the corresponding loss, effectively reducing the informativeness of the data.

[*]Equal contribution [1]Northeastern University, China [2]Lehigh University [3]Drexel University. Correspondence to: Lichao Sun <lis221@lehigh.edu>.

Proceedings of the 41st International Conference on Machine Learning, Vienna, Austria. PMLR 235, 2024. Copyright 2024 by the author(s).

Nevertheless, applying this method directly to medical images may not be optimal, as it overlooks the unique properties inherent in medical images (Liu et al., 2023d). The most important property that is overlooked is sparsity (Ye & Liu, 2012; Chuang et al., 2007; Huang et al., 2009; Otazo et al., 2015; Davoudi et al., 2019; Fang et al., 2013). For instance, in cellular microscopy images, even after cropping, there remains a substantial amount of blank background. Similarly, techniques such as CT or tomographic scanning inherently produce sparse data (Chen et al., 2018; Davoudi et al., 2019; Fang et al., 2013). Previous methods often struggled to pinpoint specific feature regions in medical data, inadvertently emphasizing sparse areas when generating noise. This led to a major waste of computational resources and suboptimal noise performance, impacts on protection effectiveness, and the time consumed for protection.

To address these challenges, we introduce a novel Sparsity-Aware Local Masking method, which leverages the inherent nature of medical data. Our approach assesses the contribution of pixels to the task based on their gradient and selects pixels with higher contributions for perturbation. This method not only narrows the perturbation search space but also enables the noise generator to focus more on feature regions, yielding noise with stronger protective effects. Additionally, since the protective performance of the noise is entirely derived from feature regions, the performance can be maintained even when cropping large background areas, a common practice in real-world medical workflows. These advancements could significantly motivate medical institutions and researchers to share their data for research or education purposes. In summary, the primary contributions of our research are as follows:

1. We are the first to find that the existing Unlearnable Example overlooks the sparse nature of medical data. Specifically, their performance is not optimal and it is not robust against medical-domain data preprocessing.

2. To address these issues, we propose Sparsity-Aware Local Masking (SALM), specifically designed for the medical domain by limiting the perturbation to features improved protection effectiveness and robustness.

3. Experiments on multiple medical datasets across different modalities, scales, and tasks demonstrate that our SALM achieves SoTA performance and is consistently effective in various medical scenarios.

2. Related Works

In this paper, we seek to protect medical data from unauthorized exploitation via a data poisoning approach. Data Poisoning is a technique used to compromise the performance of machine learning models on clean data by deliberately altering training samples. This form of attack has proven effective against both Deep Neural Networks (DNNs) and traditional machine learning methods such as SVM (Biggio et al., 2012). Muñoz-González et al. (2017) has highlighted the susceptibility of DNNs to data poisoning, although these attacks typically result in only a modest reduction in DNN performance. However, Yang et al. (2017) found there is a clear distinction between poisoned samples and normal samples, which comes at the cost of reduced data usability. Backdoor Attacks represent a specialized form of data poisoning. Traditional backdoor attacks (Chen et al., 2017; Liu et al., 2020) involve the introduction of falsely labeled training samples embedded with covert triggers into the dataset. A relatively new approach within this realm is the creation of Unlearnable Examples (Huang et al., 2021). These are considered a more subtle form of backdoor attack, free of labels and triggers. Unlearnable Examples show promising results in protecting data from unauthorized exploitation in various domains and applications (Liu et al., 2023c;b;a; Sun et al., 2022; Zhang et al., 2023; Li et al., 2023; He et al., 2022; Zhao & Lao, 2022; Salman et al., 2023; Guo et al., 2023), such as natural language processing (Ji et al., 2022), graph learning (Liu et al., 2023a), and contrastive learning (Ren et al., 2022). However, the flexibility of protecting medical data has not yet been fully explored. Recently, Liu et al. (2023d) took a first step to evaluate the performance of conventional image protection methods on medical images. Nevertheless, the methods in (Liu et al., 2023d) are straightforward adaptations from the previous approach and are suboptimal due to a lack of consideration of the intrinsic characteristics of medical data. To better address this, in this study, we leverage the inherent features of medical data to design more effective protection.

3. Methodology

3.1. Problem Statement

Assumptions on Defender's Capability. We assume that defenders are capable of making arbitrary modifications to the data they seek to protect, under the premise that these modifications do not impair the visual quality. To better simulate real-world conditions, defenders cannot interfere with the training processes of unauthorized users and do not know the specific models they use. Additionally, once the dataset is publicly released, defenders can no longer modify the data.

Objectives. This question is posed in the context of utilizing Deep Neural Networks (DNNs) for the classification of medical images. In a classification task comprising K categories, the clean training dataset and test dataset are denoted as $\mathcal{D}_c$ and $\mathcal{D}_t$, respectively.

Suppose the clean training dataset consists of n clean exam-

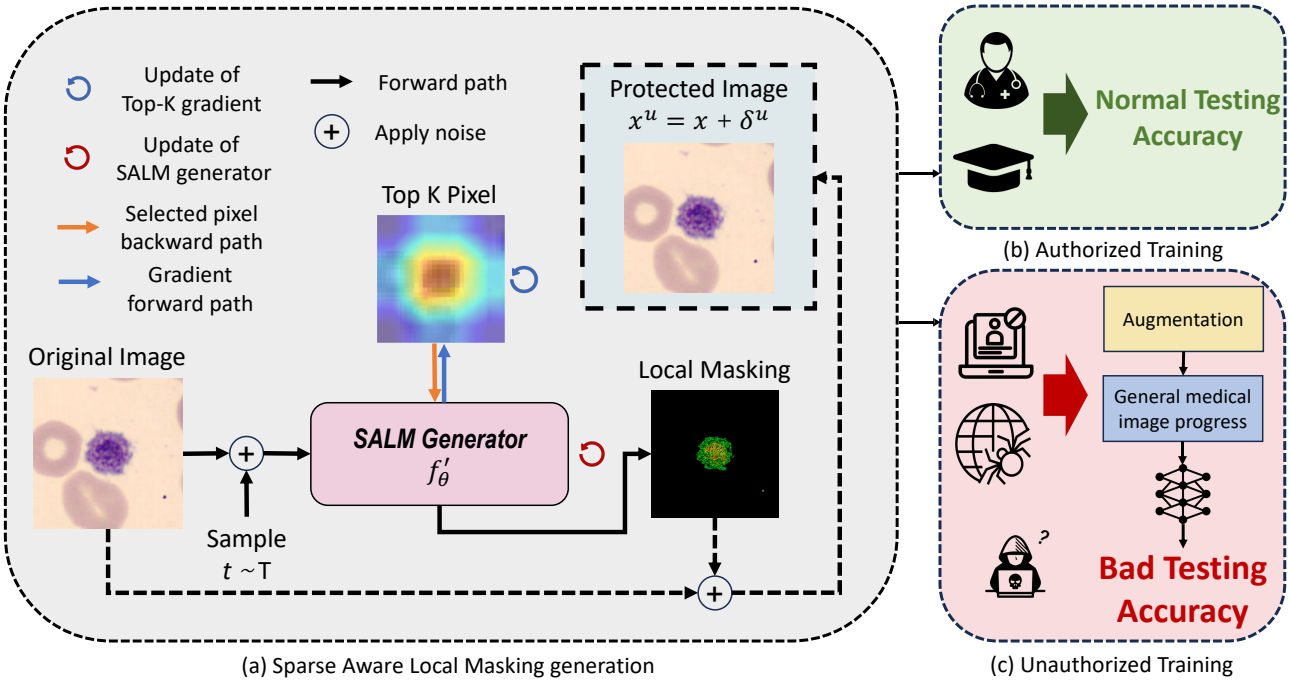

Figure 1. Our SALM method comprises a comprehensive framework that encompasses two primary steps: important pixel acquisition and noise generator training. In the first phase, the model calculates the gradient at each pixel within the image and ranks them, generating a sparse mask through a pre-set K value. In the second phase, the noise generator focuses on perturbing the pixels selected in the previous step and updates its parameters. By implementing this noise, models trained without authorization exhibit poor performance on clean datasets. Conversely, the performance for authorized users remains comparable to that achieved with the original data.

ples, that is, $\mathcal{D}_c = \{(\boldsymbol{x}_i, y_i)\}_{i=1}^n$ with $\boldsymbol{x} \in \mathcal{X} \subset \mathbb{R}^d$ are the inputs and $y \in \mathcal{Y} = \{1, \cdots, N\}$ are the labels and N is the total number of classes. We denote its unlearnable version by $\mathcal{D}_u = \{(\boldsymbol{x}'_i, y_i)\}_{i=1}^n$, where $\boldsymbol{x}' = \boldsymbol{x} + \boldsymbol{\delta}$ is the unlearnable version of training example $\boldsymbol{x} \in \mathcal{D}_c$,and $\boldsymbol{\delta} \in \Delta \subset \mathbb{R}^d$ is the "invisible" noise that makes $\boldsymbol{x}$ unlearnable. The noise $\boldsymbol{\delta}$ is bounded by $\|\delta\|_p \leq \epsilon$ with $\|\cdot\|_p (L_p$ norm), and ϵ is set to be small such that it does not affect the normal utility of the example.

In the specific scenario, a DNN model f_θ parameterized with θ, is trained on $\mathcal{D}_c$ to learn the mapping from the input domain to the output domain. For simplicity, we will omit the θ notation in the rest of this paper. To generate an unlearnable dataset, our objective is to induce the model to learn a spurious correlation between noise and labels: $f_\theta : \Delta \to \mathcal{Y}, \Delta \neq \mathcal{X}$, when trained on $\mathcal{D}_u$:

$$\min_\theta \mathbb{E}_{(\boldsymbol{x}', y) \sim \mathcal{D}_u} \mathcal{L}\left(f_\theta\left(\boldsymbol{x}'\right), y\right). \tag{1}$$

3.2. Sparsity-Aware Local Masking

In an ideal scenario, the generation of noise involves a class-matching process to protect each category within $\mathcal{D}_c$. For simplicity, we define the noise generation process on $\mathcal{D}_c$ here. Given a clean example $\boldsymbol{x}$, by generating the imperceptible noise $\boldsymbol{\Delta}$ for the training input $\boldsymbol{x}$ by solving the following bi-level optimization problem:

$$\min_\theta \mathbb{E}_{(\boldsymbol{x}, y) \sim \mathcal{D}_c} \left[\min_{\boldsymbol{\delta}} \mathcal{L}\left(f'_\theta(\boldsymbol{x} + \boldsymbol{\delta}), y\right) \right], \text{ s.t. } \|\boldsymbol{\delta}\|_p \leq \epsilon, \tag{2}$$

where f' denotes the source model used for noise generation. Note that this is a min-min bi-level optimization problem: the inner minimization is a constrained optimization problem that finds the L_p-norm bounded noise δ that minimizes the model's classification loss, while the outer minimization problem finds the parameters θ that also minimize the model's classification loss. However, such an approach does not impose any constraints on the selection of pixels for perturbation. If this method is directly applied to medical data, it overlooks the sparsity inherent in medical datasets.

To bridge the gap in the biomedical domain, motivated by the observation that the important features in the biomedical image are often sparse as we observed before, we propose a sparsity-aware objective that only modifies a portion of important pixels in the image $\boldsymbol{x}$. Formally, we introduce an additional constraint to limit the noise $\boldsymbol{\delta}$ in terms of ℓ_0 sparsity norm, i.e., $\|\boldsymbol{\delta}\|_0 \leq m$, we have

$$\min_{\theta,\boldsymbol{\delta}:\|\boldsymbol{\delta}\|_p\leq\epsilon \text{ and } \|\boldsymbol{\delta}\|_0\leq m} \mathbb{E}_{(\boldsymbol{x},y)\sim\mathcal{D}_c}\mathcal{L}(f_\theta(\boldsymbol{x}+\boldsymbol{\delta}),y). \quad (3)$$

To address the bi-level optimization problem in Eq. (2), existing studies have proposed methods such as iterative generator training (Fu et al., 2022), target poisoning with a pretrained model (Fowl et al., 2021). Specifically, in this work, we adopt the iterative generator training framework, with the training termination condition being solely the training steps M. When the training step M is sufficient, the noise will be continued optimizing to achieve better performance since there's no accuracy stop condition like (Huang et al., 2021). In each step of noise update, we employ the PGD (Madry et al., 2017) to solve the constrained minimization problem as follows:

$$\delta_{t+1}^u = \prod_{\|\delta\|_p\leq\rho_u} \left(\delta_t^u - \alpha_u \cdot \operatorname{sign}\left(\nabla_{\boldsymbol{x}}\mathcal{L}\left(f_\theta'\left(x_t+\delta_t^u\right),y\right)\right)\right), \quad (4)$$

where t is the current step in the training process, $\nabla_{\boldsymbol{x}}\mathcal{L}\left(f_\theta'\left(x_t'\right)\right)$ is the gradient of loss, Π is a projection function that clips the noise back to the refined area around the original example x when it goes beyond, α_u is the step size.

Existing methods do not account for the additional ℓ_0 constraint as proposed in Eq (3). In this paper, we leverage the principle that a pixel's contribution to the task is proportional to the magnitude of its gradient. Using the gradient as a basis, we rank pixel importance and subsequently generate a localized mask. This approach aims to achieve both high performance and efficiency. Specifically, the gradient corresponding to $\boldsymbol{x}$ is defined as $G_x = \nabla_{\boldsymbol{x}}\mathcal{L}\left(f_\theta'\left(\boldsymbol{x}_t'\right),y\right)$. The projection $\mathcal{P}$ is then applied to the obtained gradient:

$$\mathcal{P}\left(G_x,k\right) = M \odot G_x, M_{ij} = \begin{cases} 1, g_{ij} \geq g^0(k) \\ 0, \text{ otherwise} \end{cases}, \quad (5)$$

where we seek to modify the top-k percent of the pixels, g_{ij} denotes the value of position (i,j) of the gradient G_x, and $g^0(k)$ is the k-th percentile value in G_x. If $M_{ij}=0$, it means any modification at position (i,j) will result in δ_u not satisfying the constraint conditions of $\|\boldsymbol{\delta}\|_0 \leq m$.

With projection $\mathcal{P}$ to refine the local region, the perturbation δ is crafted with a PGD (Madry et al., 2017) process. Given total PGD steps K_a, for each iteration $t \in [1, K_a]$ the noise is iteratively updated via:

$$\delta_{t+1}^u = \prod_{\|\delta\|_p\leq\rho_u} \left(\delta_t^u - \alpha_u \cdot \operatorname{sign}\left(\mathcal{P}\left(g_k,k\right)_t,y\right)\right). \quad (6)$$

By purposefully reducing the range of perturbation, this approach enables the noise to cover specific feature regions rather than dispersing throughout the entire color space. This ensures that the noise does not expend effort in sparse regions unnecessarily, leading to better performance and efficiency. The overall framework and procedure are depicted in Figure 1 and Appendix B.

Table 1. Clean test accuracy (%) of RN-18 trained on datasets across various modalities protected by different methods. The symbol ↓ in the following context indicates a decrease in accuracy compared to the clean test accuracy.

Dataset	Clean	AdvT	EM	SP	SALM($K=10$)
PathMNIST	90.7	**8.4(↓82.3)**	13.5(↓77.2)	12.2(↓78.5)	11.8(↓78.9)
ChestMNIST	94.7	38.8(↓55.9)	78.7(↓16.0)	63.5(↓31.2)	**27.6(↓67.1)**
DermaMNIST	73.5	38.7(↓34.8)	17.1(↓56.4)	11.5(↓62.0)	**8.7(↓64.8)**
OCTMNIST	74.3	21.4(↓52.9)	25.0(↓49.3)	22.7(↓51.5)	**21.2(↓53.1)**
PneumoniaMNIST	85.4	66.9(↓18.5)	65.8(↓19.6)	70.1(↓15.3)	**63.9(↓21.5)**
RetinaMNIST	52.4	39.5(↓12.9)	13.0(↓39.4)	**8.3(↓44.1)**	12.2(↓40.2)
BreastMNIST	86.3	**44.9(↓41.4)**	46.2(↓40.1)	50.0(↓36.3)	53.2(↓33.1)
BloodMNIST	95.8	19.6(↓76.2)	30.5(↓65.3)	30.5(↓65.3)	**7.1(↓88.7)**
TissueMNIST	67.6	19.7(↓47.9)	17.4(↓50.2)	**7.3(↓60.3)**	21.4(↓46.2)
OrganAMNIST	93.5	81.0(↓12.5)	78.1(↓15.4)	86.1(↓7.4)	**69.0(↓24.5)**
OrganCMNIST	90.0	70.3(↓19.7)	72.2(↓17.8)	74.8(↓15.2)	**64.6(↓25.4)**
OrganSMNIST	78.2	51.7(↓26.5)	50.1(↓28.1)	53.0(↓25.2)	**46.8(↓31.4)**
PathMNIST_{224}	90.9	70.8(↓21.9)	32.5(↓60.2)	29.1(↓63.6)	**27.5(↓67.3)**
ChestMNIST_{224}	92.7	31.6(↓61.1)	64.4(↓28.3)	50.1(↓42.6)	**25.4(↓67.3)**

4. Experiments and Results

We selected more than 14 datasets from Medmnist (Yang et al., 2021) and Medmnist-v2 (Yang et al., 2023) and conducted extensive experiments on various models following (Huang et al., 2021). We compare SALM with three baselines, including **AdvT** (Fowl et al., 2021), **EM** (Huang et al., 2021), **SP** (Yu et al., 2022). For more experimental details, please refer to Appendix C. More experiments, analysis, and the case study can be found in Appendix D and E.

Effectiveness Analysis. We initially selected PathMNIST for comparison between our method and the conceptually similar error-minimizing and error-maximizing noise. The results in Figure 2 show that from the onset of training, both our method and the error-minimizing noise provide effective protection, with our method demonstrating superior performance. Conversely, the efficacy of error-maximizing noise gradually diminished throughout the training. To further evaluate the performance of our method, we directly transferred several conventional image protection methods to medical and compared them with our approach. The results in Table 1 confirm that our method is not only broadly applicable to different modalities of medical datasets but also outperforms the previous SoTA methods.

Different Architectures Selection. Before releasing datasets, we cannot foresee the specific training model unauthorized users might adopt. Hence, ensuring that the protected data remains effective across various models is crucial. While protectors can only select a single source model for noise generation, unauthorized users are free to choose any model, seemingly placing the protectors at a disadvantage place. However, in reality, the choice of model

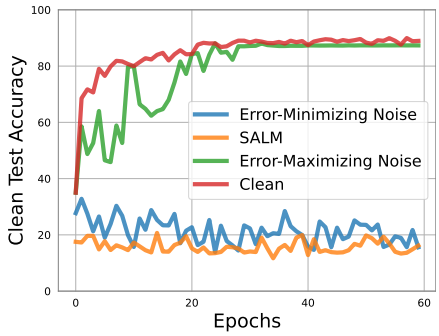

Figure 2. The learning curves of ResNet-18 trained on different protected data.

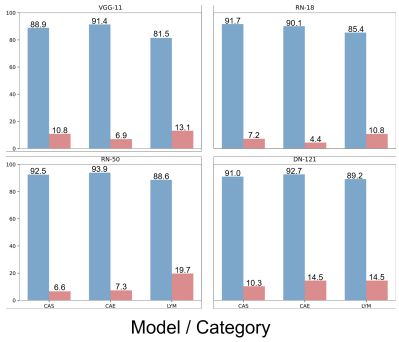

Figure 3. The selected categories protect effectiveness under different models.

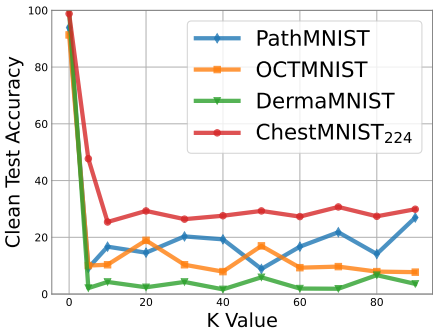

Figure 4. The effect of K on clean test accuracy(%) for the four datasets.

Table 2. Clean test accuracy (%) of DNNs trained on the clean training sets ($\mathcal{D}_c$) or their unlearnable ones ($\mathcal{D}_u$) made by different K value.

K	Model	PathMNIST		DermaMNIST		OctMNIST		ChestMNIST$_{224}$	
		$\mathcal{D}_c$	$\mathcal{D}_u$	$\mathcal{D}_c$	$\mathcal{D}_u$	$\mathcal{D}_c$	$\mathcal{D}_u$	$\mathcal{D}_c$	$\mathcal{D}_u$
$K = 10$	VGG-11	91.7	**19.2**	73.9	**10.9**	78.2	**21.4**	91.7	**27.3**
	RN-18	90.7	**11.8**	73.5	**8.7**	74.3	**21.2**	92.7	**25.4**
	RN-50	94.6	**21.9**	74.6	**9.8**	74.3	**19.0**	94.7	**25.4**
	DN-121	96.4	**14.8**	75.3	**14.5**	75.6	**13.3**	95.1	**21.2**
$K = 30$	VGG-11	90.8	**18.3**	74.4	**11.8**	75.5	**15.1**	92.1	**31.4**
	RN-18	92.9	**14.1**	73.5	**4.0**	74.4	**12.2**	94.6	**24.0**
	RN-50	94.2	**17.4**	74.4	**9.3**	77.0	**14.3**	94.8	**30.7**
	DN-121	94.2	**19.7**	77.2	**11.8**	74.9	**30.0**	94.5	**29.9**

Table 3. Clean test accuracy (%) of RN-18 trained on datasets protected by different methods after three common low-pass filters.

Dataset	Filter	Clean	AdvT	EM	SP	SALM
PathMNIST	Mean	91.3	17.9(↓73.4)	16.4(↓74.9)	17.1(↓74.2)	**15.6(↓75.7)**
	Median	90.8	22.4(↓68.4)	**7.2(↓83.6)**	12.7(↓78.1)	19.3(↓71.5)
	Gaussian	92.1	35.5(↓56.6)	**11.0(↓81.1)**	18.0(↓74.1)	17.7(↓74.4)
DermaMNIST	Mean	73.8	45.8(↓28.0)	31.9(↓41.9)	23.9(↓49.9)	**21.5(↓52.3)**
	Median	75.1	62.6(↓12.5)	34.6(↓40.5)	19.9(↓55.2)	**17.8(↓57.3)**
	Gaussian	74.7	40.1(↓34.6)	33.1(↓41.6)	**25.8(↓48.9)**	40.6(↓34.1)
OctMNIST	Mean	77.7	25.7(↓52.0)	28.3(↓49.4)	**23.8(↓53.9)**	29.6(↓48.1)
	Median	79.6	23.5(↓56.1)	25.7(↓53.9)	23.6(↓56.0)	**22.6(↓57.0)**
	Gaussian	77.7	**22.9(↓54.8)**	26.1(↓51.6)	23.1(↓54.6)	29.2(↓48.5)
ChestMNIST$_{224}$	Mean	91.4	37.7(↓53.7)	72.5(↓18.9)	58.2(↓33.2)	**27.0(↓64.4)**
	Median	90.9	38.2(↓52.7)	62.7(↓28.2)	50.9(↓40.0)	**28.4(↓62.5)**
	Gaussian	91.1	45.1(↓46.0)	70.8(↓20.3)	51.4(↓39.7)	**29.8(↓61.3)**

by unauthorized users does not impact the efficacy of our method. The results in Table 2 demonstrate that the SALM with the source model ResNet-18 (He et al., 2016) is effective across various models. The results in Figure 3 show the randomly extracted subset still has excellent performance under different models. Furthermore, the protection performance is not necessarily best when unauthorized trained on ResNet-18. This indicates that our method is not limited by the source model and the target model, enabling more effective application in real-world scenarios and facilitating the open-sourcing of high-quality datasets. What's more, the protective effect when K is 30 is not significantly better than when it is 10. This prompts us to delve deeper into the impact of the choice of K-value.

Different K Selection. In this section, we delve deeper into our approach to unveil the nuances of our method's dependency on the choice of K value. K values were selected within a range from 0 to 90 in increments of ten. We also tried $K = 5$ to test the limitation of SALM. When K is set to 0, it implies that the model is trained on clean data. The results in Figure 4 show when K is set at 5 or higher, the noise proves to be generally effective. Consequently, in subsequent sections, unless specifically emphasized, we opt for $K = 10$ for our experiments. This underscores the effectiveness of limiting the perturbation search space and also demonstrates that the existing methods are not optimal in the medical domain. What's more, it highlights the high efficiency of our approach to medical data, which can enable the accelerated release of datasets.

Resistance to Low-pass Filters. Constrained by limitations in the quality of medical imaging, the acquired images often exhibit noise. Consequently, researchers commonly employ low-pass filters for pre-processing due to their simplicity and efficiency. Therefore, verifying the resilience of our method to such filtering is of paramount importance. We chose three low-pass filters: Mean, Median, and Gaussian, each with a 3×3 window size. The results in Table 3 show that our method retains its effectiveness post-filtering, and the sensibility is lower than other methods, affirming its applicability within actual medical workflows.

5. Conclusion

This work introduces the SALM method, a novel approach dedicated to generating Unlearnable Examples specifically designed for medical datasets. Extensive observation has revealed that medical datasets are inherently sparse, a characteristic not effectively utilized by existing methods for generating Unlearnable Examples. Consequently, the SALM method is designed to selectively perturb only a specific subset of critical pixels. Extensive experimental results demonstrate that the SALM method effectively protects medical images from unauthorized training. Simultaneously,

it ensures stability and effectiveness throughout common medical image processing workflows (e.g., filtering and cropping). Additionally, the processed images retain their visual usability, not impeding clinical diagnosis by physicians. Furthermore, we demonstrate that SALM has strong, flexible applicability in practical application scenarios.

References

Akinyele, J. A., Pagano, M. W., Green, M. D., Lehmann, C. U., Peterson, Z. N., and Rubin, A. D. Securing electronic medical records using attribute-based encryption on mobile devices. In *Proceedings of the 1st ACM workshop on Security and privacy in smartphones and mobile devices*, pp. 75–86, 2011.

Alberdi, A., Aztiria, A., and Basarab, A. On the early diagnosis of alzheimer's disease from multimodal signals: A survey. *Artificial intelligence in medicine*, 71:1–29, 2016.

Alpert, S. A. Protecting medical privacy: challenges in the age of genetic information. *Journal of Social Issues*, 59 (2):301–322, 2003.

Baowaly, M. K., Lin, C.-C., Liu, C.-L., and Chen, K.-T. Synthesizing electronic health records using improved generative adversarial networks. *Journal of the American Medical Informatics Association*, 26(3):228–241, 2019.

Barrows Jr, R. C. and Clayton, P. D. Privacy, confidentiality, and electronic medical records. *Journal of the American medical informatics association*, 3(2):139–148, 1996.

Biggio, B., Nelson, B., and Laskov, P. Poisoning attacks against support vector machines. *arXiv preprint arXiv:1206.6389*, 2012.

Bilic, P., Christ, P., Li, H. B., Vorontsov, E., Ben-Cohen, A., Kaissis, G., Szeskin, A., Jacobs, C., Mamani, G. E. H., Chartrand, G., et al. The liver tumor segmentation benchmark (lits). *Medical Image Analysis*, 84:102680, 2023.

Chen, H., Zhang, Y., Chen, Y., Zhang, J., Zhang, W., Sun, H., Lv, Y., Liao, P., Zhou, J., and Wang, G. Learn: Learned experts' assessment-based reconstruction network for sparse-data ct. *IEEE transactions on medical imaging*, 37(6):1333–1347, 2018.

Chen, X., Liu, C., Li, B., Lu, K., and Song, D. Targeted backdoor attacks on deep learning systems using data poisoning. *arXiv preprint arXiv:1712.05526*, 2017.

Chuang, H.-Y., Lee, E., Liu, Y.-T., Lee, D., and Ideker, T. Network-based classification of breast cancer metastasis. *Molecular systems biology*, 3(1):140, 2007.

Davoudi, N., Deán-Ben, X. L., and Razansky, D. Deep learning optoacoustic tomography with sparse data. *Nature Machine Intelligence*, 1(10):453–460, 2019.

Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pp. 248–255. Ieee, 2009.

Edwards, B. Artist finds private medical record photos in popular ai training data set, 2022.

Fang, L., Li, S., McNabb, R. P., Nie, Q., Kuo, A. N., Toth, C. A., Izatt, J. A., and Farsiu, S. Fast acquisition and reconstruction of optical coherence tomography images via sparse representation. *IEEE transactions on medical imaging*, 32(11):2034–2049, 2013.

Forghani, R., Savadjiev, P., Chatterjee, A., Muthukrishnan, N., Reinhold, C., and Forghani, B. Radiomics and artificial intelligence for biomarker and prediction model development in oncology. *Computational and structural biotechnology journal*, 17:995, 2019.

Fowl, L., Goldblum, M., Chiang, P.-y., Geiping, J., Czaja, W., and Goldstein, T. Adversarial examples make strong poisons. *Advances in Neural Information Processing Systems*, 34:30339–30351, 2021.

Fu, S., He, F., Liu, Y., Shen, L., and Tao, D. Robust unlearnable examples: Protecting data against adversarial learning. *arXiv preprint arXiv:2203.14533*, 2022.

Gong, T., Huang, H., Li, P., Zhang, K., and Jiang, H. A medical healthcare system for privacy protection based on iot. In *2015 Seventh International Symposium on Parallel Architectures, Algorithms and Programming (PAAP)*, pp. 217–222. IEEE, 2015.

Gostin, L. O. *Public health law: power, duty, restraint*, volume 3. Univ of California Press, 2000.

Gunraj, H., Wang, L., and Wong, A. Covidnet-ct: A tailored deep convolutional neural network design for detection of covid-19 cases from chest ct images. *Frontiers in medicine*, 7:608525, 2020.

Guo, J., Li, Y., Wang, L., Xia, S.-T., Huang, H., Liu, C., and Li, B. Domain watermark: Effective and harmless dataset copyright protection is closed at hand. *arXiv preprint arXiv:2310.14942*, 2023.

He, H., Zha, K., and Katabi, D. Indiscriminate poisoning attacks on unsupervised contrastive learning. *arXiv preprint arXiv:2202.11202*, 2022.

He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE*

conference on computer vision and pattern recognition, pp. 770–778, 2016.

Heurix, J. and Neubauer, T. Privacy-preserving storage and access of medical data through pseudonymization and encryption. In *Trust, Privacy and Security in Digital Business: 8th International Conference, TrustBus 2011, Toulouse, France, August 29-September 2, 2011. Proceedings 8*, pp. 186–197. Springer, 2011.

Hill, K. and Krolik, A. How photos of your kids are powering surveillance technology. *The New York Times*, 2019.

Huang, G., Liu, Z., Van Der Maaten, L., and Weinberger, K. Q. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4700–4708, 2017.

Huang, H., Ma, X., Erfani, S. M., Bailey, J., and Wang, Y. Unlearnable examples: Making personal data unexploitable. *arXiv preprint arXiv:2101.04898*, 2021.

Huang, S., Li, J., Sun, L., Liu, J., Wu, T., Chen, K., Fleisher, A., Reiman, E., and Ye, J. Learning brain connectivity of alzheimer's disease from neuroimaging data. *Advances in Neural Information Processing Systems*, 22, 2009.

Irvin, J., Rajpurkar, P., Ko, M., Yu, Y., Ciurea-Ilcus, S., Chute, C., Marklund, H., Haghgoo, B., Ball, R., Shpanskaya, K., et al. Chexpert: A large chest radiograph dataset with uncertainty labels and expert comparison. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pp. 590–597, 2019.

Ji, Z., Ma, P., and Wang, S. Unlearnable examples: Protecting open-source software from unauthorized neural code learning. In *SEKE*, pp. 525–530, 2022.

Johnson, A. E., Pollard, T. J., Greenbaum, N. R., Lungren, M. P., Deng, C.-y., Peng, Y., Lu, Z., Mark, R. G., Berkowitz, S. J., and Horng, S. Mimic-cxr-jpg, a large publicly available database of labeled chest radiographs. *arXiv preprint arXiv:1901.07042*, 2019.

Kather, J. N., Krisam, J., Charoentong, P., Luedde, T., Herpel, E., Weis, C.-A., Gaiser, T., Marx, A., Valous, N. A., Ferber, D., et al. Predicting survival from colorectal cancer histology slides using deep learning: A retrospective multicenter study. *PLoS medicine*, 16(1):e1002730, 2019.

Kayaalp, M. Patient privacy in the era of big data. *Balkan medical journal*, 35(1):8–17, 2018.

Kelly, C. J., Karthikesalingam, A., Suleyman, M., Corrado, G., and King, D. Key challenges for delivering clinical impact with artificial intelligence. *BMC medicine*, 17: 1–9, 2019.

Kermany, D. S., Goldbaum, M., Cai, W., Valentim, C. C., Liang, H., Baxter, S. L., McKeown, A., Yang, G., Wu, X., Yan, F., et al. Identifying medical diagnoses and treatable diseases by image-based deep learning. *cell*, 172(5):1122–1131, 2018.

Koh, H. C., Tan, G., et al. Data mining applications in healthcare. *Journal of healthcare information management*, 19 (2):65, 2011.

Li, C.-T., Lee, C.-C., and Weng, C.-Y. A secure cloud-assisted wireless body area network in mobile emergency medical care system. *Journal of medical systems*, 40: 1–15, 2016.

Li, Z., Yu, N., Salem, A., Backes, M., Fritz, M., and Zhang, Y. {UnGANable}: Defending against {GAN-based} face manipulation. In *32nd USENIX Security Symposium (USENIX Security 23)*, pp. 7213–7230, 2023.

Liu, C.-H., Chung, Y.-F., Chen, T.-S., and Wang, S.-D. The enhancement of security in healthcare information systems. *Journal of medical systems*, 36:1673–1688, 2012.

Liu, F. and Li, T. A clustering k-anonymity privacy-preserving method for wearable iot devices. *Security and Communication Networks*, 2018:1–8, 2018.

Liu, Y., Ma, X., Bailey, J., and Lu, F. Reflection backdoor: A natural backdoor attack on deep neural networks. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part X 16*, pp. 182–199. Springer, 2020.

Liu, Y., Fan, C., Chen, X., Zhou, P., and Sun, L. Graphcloak: Safeguarding task-specific knowledge within graph-structured data from unauthorized exploitation. *arXiv preprint arXiv:2310.07100*, 2023a.

Liu, Y., Fan, C., Dai, Y., Chen, X., Zhou, P., and Sun, L. Toward robust imperceptible perturbation against unauthorized text-to-image diffusion-based synthesis. *arXiv preprint arXiv:2311.13127*, 2023b.

Liu, Y., Fan, C., Zhou, P., and Sun, L. Unlearnable graph: Protecting graphs from unauthorized exploitation. *arXiv preprint arXiv:2303.02568*, 2023c.

Liu, Y., Ye, H., Zhang, K., and Sun, L. Securing biomedical images from unauthorized training with anti-learning perturbation. *arXiv preprint arXiv:2303.02559*, 2023d.

Madry, A., Makelov, A., Schmidt, L., Tsipras, D., and Vladu, A. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017.

Muñoz-González, L., Biggio, B., Demontis, A., Paudice, A., Wongrassamee, V., Lupu, E. C., and Roli, F. Towards poisoning of deep learning algorithms with back-gradient optimization. In *Proceedings of the 10th ACM workshop on artificial intelligence and security*, pp. 27–38, 2017.

Murdoch, B. Privacy and artificial intelligence: challenges for protecting health information in a new era. *BMC Medical Ethics*, 22(1):1–5, 2021.

Otazo, R., Candes, E., and Sodickson, D. K. Low-rank plus sparse matrix decomposition for accelerated dynamic mri with separation of background and dynamic components. *Magnetic resonance in medicine*, 73(3):1125–1136, 2015.

Pedrosa, J., Aresta, G., Ferreira, C., Rodrigues, M., Leitão, P., Carvalho, A. S., Rebelo, J., Negrão, E., Ramos, I., Cunha, A., et al. Lndb: a lung nodule database on computed tomography. *arXiv preprint arXiv:1911.08434*, 2019.

Price, W. N. and Cohen, I. G. Privacy in the age of medical big data. *Nature medicine*, 25(1):37–43, 2019.

Rajkomar, A., Oren, E., Chen, K., Dai, A. M., Hajaj, N., Hardt, M., Liu, P. J., Liu, X., Marcus, J., Sun, M., et al. Scalable and accurate deep learning with electronic health records. *NPJ digital medicine*, 1(1):18, 2018.

Rajpurkar, P., Irvin, J., Bagul, A., Ding, D., Duan, T., Mehta, H., Yang, B., Zhu, K., Laird, D., Ball, R. L., et al. Mura: Large dataset for abnormality detection in musculoskeletal radiographs. *arXiv preprint arXiv:1712.06957*, 2017.

Rasmy, L., Xiang, Y., Xie, Z., Tao, C., and Zhi, D. Medbert: pretrained contextualized embeddings on large-scale structured electronic health records for disease prediction. *NPJ digital medicine*, 4(1):86, 2021.

Ren, J., Xu, H., Wan, Y., Ma, X., Sun, L., and Tang, J. Transferable unlearnable examples. *arXiv preprint arXiv:2210.10114*, 2022.

Salman, H., Khaddaj, A., Leclerc, G., Ilyas, A., and Madry, A. Raising the cost of malicious ai-powered image editing. *arXiv preprint arXiv:2302.06588*, 2023.

Senbekov, M., Saliev, T., Bukeyeva, Z., Almabayeva, A., Zhanaliyeva, M., Aitenova, N., Toishibekov, Y., Fakhradiyev, I., et al. The recent progress and applications of digital technologies in healthcare: a review. *International journal of telemedicine and applications*, 2020, 2020.

Shan, S., Wenger, E., Zhang, J., Li, H., Zheng, H., and Zhao, B. Y. Fawkes: Protecting privacy against unauthorized deep learning models. In *29th USENIX security symposium (USENIX Security 20)*, pp. 1589–1604, 2020.

Simonyan, K. and Zisserman, A. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.

Simpson, A. L., Antonelli, M., Bakas, S., Bilello, M., Farahani, K., Van Ginneken, B., Kopp-Schneider, A., Landman, B. A., Litjens, G., Menze, B., et al. A large annotated medical image dataset for the development and evaluation of segmentation algorithms. *arXiv preprint arXiv:1902.09063*, 2019.

Sun, J., Yao, X., Wang, S., and Wu, Y. Blockchain-based secure storage and access scheme for electronic medical records in ipfs. *IEEE access*, 8:59389–59401, 2020.

Sun, Z., Du, X., Song, F., Ni, M., and Li, L. Coprotector: Protect open-source code against unauthorized training usage with data poisoning. In *Proceedings of the ACM Web Conference 2022*, pp. 652–660, 2022.

Trinidad, M. G., Platt, J., and Kardia, S. L. The public's comfort with sharing health data with third-party commercial companies. *Humanities and Social Sciences Communications*, 7(1):1–10, 2020.

Yang, C., Wu, Q., Li, H., and Chen, Y. Generative poisoning attack method against neural networks. *arXiv preprint arXiv:1703.01340*, 2017.

Yang, J., Shi, R., and Ni, B. Medmnist classification decathlon: A lightweight automl benchmark for medical image analysis. In *IEEE 18th International Symposium on Biomedical Imaging (ISBI)*, pp. 191–195, 2021.

Yang, J., Shi, R., Wei, D., Liu, Z., Zhao, L., Ke, B., Pfister, H., and Ni, B. Medmnist v2-a large-scale lightweight benchmark for 2d and 3d biomedical image classification. *Scientific Data*, 10(1):41, 2023.

Ye, J. and Liu, J. Sparse methods for biomedical data. *ACM Sigkdd Explorations Newsletter*, 14(1):4–15, 2012.

Yoon, J., Drumright, L. N., and Van Der Schaar, M. Anonymization through data synthesis using generative adversarial networks (ads-gan). *IEEE journal of biomedical and health informatics*, 24(8):2378–2388, 2020.

Yu, D., Zhang, H., Chen, W., Yin, J., and Liu, T.-Y. Availability attacks create shortcuts. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 2367–2376, 2022.

Zhang, A., Xing, L., Zou, J., and Wu, J. C. Shifting machine learning for healthcare from development to deployment and from models to data. *Nature Biomedical Engineering*, 6(12):1330–1345, 2022.

Zhang, J., Ma, X., Yi, Q., Sang, J., Jiang, Y.-G., Wang, Y., and Xu, C. Unlearnable clusters: Towards label-agnostic unlearnable examples. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3984–3993, 2023.

Zhang, Y., Ling, H., Gao, J., Yin, K., Lafleche, J.-F., Barriuso, A., Torralba, A., and Fidler, S. Datasetgan: Efficient labeled data factory with minimal human effort. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10145–10155, 2021.

Zhao, B. and Lao, Y. Clpa: Clean-label poisoning availability attacks using generative adversarial nets. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pp. 9162–9170, 2022.

A. More Related Work

Medical Data Protection. As information technology advances, digital technologies are increasingly integrated into medicine, affecting both clinical treatment and scientific research (Senbekov et al., 2020; Rajpurkar et al., 2017; Johnson et al., 2019; Irvin et al., 2019; Pedrosa et al., 2019; Bilic et al., 2023; Rasmy et al., 2021; Zhang et al., 2022; Kelly et al., 2019). However, these developments also present significant risks and challenges regarding patient privacy, medical information breaches occur frequently around the world (Alpert, 2003; Kayaalp, 2018; Price & Cohen, 2019; Murdoch, 2021). Inadequate data protection can lead to substantial harm through the leakage of personal information, such as the disclosure of sensitive health conditions like HIV, potentially leading to social isolation and psychological disorders (Gostin, 2000). Furthermore, incidents of data breaches may also reduce patients' trust in medical research institutions, making them reluctant to share their data (Koh et al., 2011; Trinidad et al., 2020). Earlier studies focused on robust physical protection measures, like using encrypted storage devices (Heurix & Neubauer, 2011; Akinyele et al., 2011; Sun et al., 2020), establishing firewalls (Barrows Jr & Clayton, 1996; Liu et al., 2012), and secure communication transmission modes (Gong et al., 2015; Li et al., 2016), to safeguard data. Liu & Li (2018) introduced a clustering method based on K-anonymity algorithm as the building block of privacy-preserving for medical devices' data. However, as collaboration among various research institutions intensifies and open-source data sharing on the internet becomes more crucial, these methods become less applicable. Baowaly et al. (2019); Yoon et al. (2020) attempted to use generated data to reduce granularity and thus protect privacy, but this method involves a significant trade-off between information loss and protection efficacy. However, our method remains robust under various conditions during the construction process of real-world datasets and effectively balances high data usability with consistent protection.

B. SALM Algorithm

Algorithm 1 Training SALM generator and generating noise.

Input: Training data set $\mathcal{T}$, Training steps M, PGD parameters α_u and ρ_u, transformation distribution T, the sampling number J for gradient approximation.
Initialization: source model parameter θ, δ^u.
// Following Eq. (3)
for i **in** $1, \cdots, M$ **do**
  Sample minibatch $(x, y) \sim \mathcal{T}$, sample transformation $t_j \sim T$
  Calculate $g_k \leftarrow \frac{1}{J}\sum_{j=1}^{J} \frac{\partial}{\partial \delta^u} \ell(f'_\theta(t_j(x+\delta^u)), y)$
  Determine k-th percentile value $g^0(k)$ in g_k
  for each element g_{ij} in g_k **do**
    if $g_{ij} \geq g^0(k)$ **then**
      Set $M_{ij} = 1$
    else
      Set $M_{ij} = 0$
    end if
  end for
  Apply local mask: $\mathcal{P}(g_k, k) = M \odot g_k$ // Following Eq. (5)
  Update $\delta^u \leftarrow \prod_{\|\delta\| \leq \rho_u} (\delta^u - \alpha_u \cdot \text{sign}(\mathcal{P}(g_k, k)))$
  Update source model parameter θ based on minibatch $(t(x+\delta^u), y)$
end for
Output: SALM noise generator f'_θ, SALM noise δ^u

C. Experiments Setup

Datasets. Our research involved extensive experiments on MedMNIST (Yang et al., 2021), which is a comprehensive collection of standardized biomedical images spanning 12 distinct datasets. MedMNIST includes the major modalities in medical imaging. This dataset showcases diverse data scales and a range of tasks. Furthermore, to more effectively assess our method's performance on large, real-world datasets, we also do experiments on the 224x224 datasets from MedMNIST-v2 (Yang et al., 2023).

Table 4. More results of RN-18 trained on 224 version datasets protected by different methods. The symbol ↓ in the following context indicates a decrease in accuracy compared to the clean test accuracy.

Dataset	Clean	AdvT	EM	SP	SALM($K = 10$)
DermaMNIST_{224}	74.8	19.2(↓55.6)	11.4(↓63.4)	10.0(↓64.8)	**9.3(↓65.5)**
OCTMNIST_{224}	76.2	32.3(↓43.9)	**14.6(↓61.6)**	20.6(↓55.6)	25.8(↓50.4)
$\text{PneumoniaMNIST}_{224}$	84.1	66.7(↓17.4)	67.2(↓16.9)	61.6(↓22.5)	**57.4(↓26.7)**
RetinaMNIST_{224}	51.0	27.6(↓23.4)	21.1(↓29.9)	10.5(↓40.5)	**9.1(↓41.9)**
BreastMNIST_{224}	80.1	**35.1(↓45.0)**	40.0(↓40.1)	40.2(↓39.9)	37.8(↓42.3)
BloodMNIST_{224}	96.0	16.3(↓80.3)	16.6(↓79.4)	7.0(↓89.0)	**6.3(↓89.7)**
TissueMNIST_{224}	68.4	18.5(↓49.9)	23.6(↓44.8)	**16.7(↓51.7)**	19.1(↓49.3)
OrganAMNIST_{224}	95.2	85.8(↓9.4)	83.5(↓11.7)	85.3(↓9.9)	**82.1(↓13.1)**
OrganCMNIST_{224}	91.6	82.3(↓9.3)	69.7(↓21.9)	72.2(↓19.4)	**69.3(↓22.3)**
OrganSMNIST_{224}	77.8	47.8(↓30.0)	44.6(↓33.2)	65.2(↓12.6)	**42.1(↓35.7)**

Table 5. Test accuracy (%) of RN-18 trained on different kinds of data of PathMNIST.

Method	$\mathcal{D}_c$	Cropped($\mathcal{D}_c$)	$\mathcal{D}_u$	Cropped($\mathcal{D}_u$)	Gap
TAP	91.2	81.9	26.5	12.7	4.5 ↓
EM	90.6	82.2	17.2	4.7	4.1 ↓
SP	90.9	81.6	21.3	12.3	0.2 ↑
SALM	90.9	82.5	14.9	5.3	**1.2**↑

Table 6. Image similarity scores of different methods.

Method	SSIM	aHash	pHash	dHash	NMI
TAP	94.6	75.1	88.8	77.2	42.7
EM	92.7	77.7	89.1	77.4	42.9
SP	99.9	95.4	99.9	99.3	99.9
SALM	94.4	77.9	89.4	77.2	44.7

Model and Implementation Details. We selected the following well-known models to evaluate our method against benchmark approaches: VGG-11 (Simonyan & Zisserman, 2014), ResNet-18/50 (He et al., 2016) and DenseNet-121 (Huang et al., 2017). For the SALM generator, we select RN-18 as the source model to generate noise. We set the perturbation noise $\rho_u = 8/255$ as default. We choose an SGD optimizer for both noise generation and training. Both of their weight decay are set to 5e-4, and momentum is set to 0.9. The initial learning rate is set to 0.1 with a decay rate of 0.1.

Baselines. Since there's no previous method specifically designed for medical data, we compare our Sparsity-Aware Local Masking **(SALM)** with the existing SoTA methods in the general image domain. The baseline methods are as following: **T**arget **A**dversarial **P**oisoning **(TAP)** (Fowl et al., 2021), **E**rror-**M**inimizing noise **(EM)** (Huang et al., 2021), **S**ynthesized **P**erturbation **(SP)** (Yu et al., 2022). None of these three methods impose restrictions on the range of pixels to be perturbed. TAP and EM both target the training loss and craft noise using the gradient information to trigger maximum error or trick models to overfit. SC, on the other hand, imposes hand-crafted linear-separable noise into the data, thereby leading the model to learn only simple noise-label correlation.

D. More Experiments

D.1. Resistance to Cropping.

As we consistently highlighted, medical data exhibits greater sparsity compared to conventional data. For instance, when processing microscopic images, physicians or researchers typically reduce the proportion of the background through cropping and magnifying details to facilitate analysis. In this section, we further validate the effectiveness of our SALM after cropping. We calculate the difference in test accuracy on the protected data before and after cropping. To minimize the impact of cropping itself on performance, we add the difference in test accuracy on the clean data before and after cropping, which we denote as the Gap value. The results in Table 5 show after cropping, the performance gap of TAP and EM significantly narrows. This finding confirms that color space disturbances by these methods are random. If perturbations are concentrated in non-feature areas like transparent slides, cropping might result in a loss of protection. Conversely, SALM focuses on pixels that significantly contribute to the model, which ensures that removing the background does not impact the regions with more protective noise and thus preserves the protection.

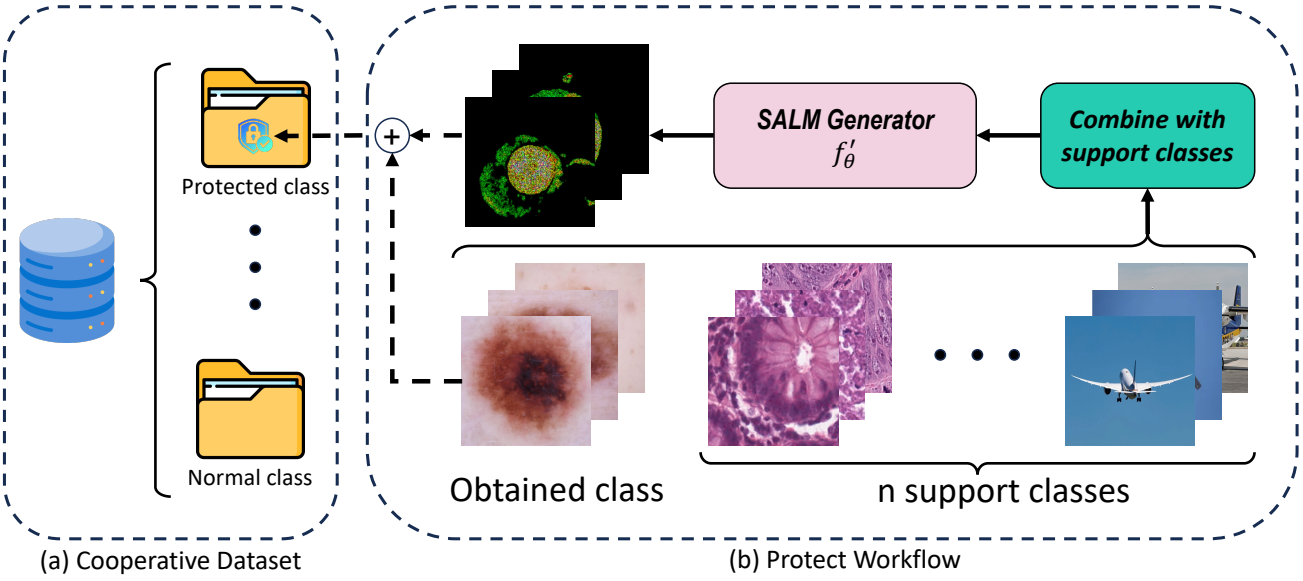

Figure 5. The protected framework in the case of the class-combined dataset. The noise generated for the corresponding obtained class from data combined with other support classes still retains its protective effect.

D.2. Similarity Compare.

Although previous studies (Liu et al., 2023d; Huang et al., 2021; Fu et al., 2022; Fowl et al., 2021; Yu et al., 2022) have demonstrated that perturbation radii as small as 8 or even 16 do not affect human visual perception, considering the sensitivity of medical images, we still assessed the similarity. Specifically, we utilized the Structural Similarity Index Measure (SSIM), which assesses luminance, contrast, and structure, along with various Hash methods focusing on low-frequency information. The results in Table 6 reveal that SYN has the highest scores. This is likely due to its structure, consisting of square regions of differing colors with minimal distinctions within these areas. However, compared with other methods, despite being based on feature perturbations, our method does not substantially affect images. This underscores our approach's ability to effectively preserve image utility.

E. Case Study

E.1. Case 1: Class-wise Combined Dataset

We initially conducted a study on the scenario of the class-wise combination. To better explain the experimental setting, we introduce the concept of "support classes," which is utilized to augment the existing single-class data. We selected two kinds of support classes: medical classes and general classes. The framework is shown in Figure 5. First, we chose melanoma from the DermaMNIST as a protected target and selected all classes within the PathMNIST as support classes for our experiments. The melanoma data was combined with PathMNIST to create a ten-class dataset while only preserving the SALM noise generated for the melanoma. Subsequently, to broaden the selection range of support classes, we also selected some general image classes (e.g., airplanes) as support classes.

Following the generation of class-specific noise, we utilized the same experimental settings as in previous tests. Using PathMNIST for support classes, the generated noise resulted in a clean test accuracy of **9.9%**. Noise generated with general image classes from ImageNet (Deng et al., 2009) as support classes led to an accuracy of **8.1%**. The choice of either type of support class does not affect the noise performance. This demonstrates the flexible applicability of our method.

E.2. Case 2: Sample-wise Combined Dataset

In addition to the scenarios mentioned above, researchers may only be responsible for a portion of the data within the whole dataset. They extensively collect various classes of data and annotate them, then merge them into a complete dataset based

Table 7. Effectiveness under different unlearnable percentages on PathMNIST with RN-18 model: lower clean accuracy indicates better effectiveness. $\mathcal{D}_u + \mathcal{D}_c$: a mix of unlearnable and clean data; $\mathcal{D}_c$: only the clean proportion of data. Percentage of unlearnable examples: $\frac{\mathcal{D}_u}{\mathcal{D}_c+\mathcal{D}_u}$.

K Value	Percentage of unlearnable examples									
	0%	**20%**		**40%**		**60%**		**80%**		**100%**
		$\mathcal{D}_u + \mathcal{D}_c$	$\mathcal{D}_c$	$\mathcal{D}_u + \mathcal{D}_c$	$\mathcal{D}_c$	$\mathcal{D}_u + \mathcal{D}_c$	$\mathcal{D}_c$	$\mathcal{D}_u + \mathcal{D}_c$	$\mathcal{D}_c$	
$K = 10$	91.2	90.8	91.0	90.0	89.2	89.9	87.2	88.3	89.2	11.8
$K = 30$	90.9	91.3	91.2	90.2	90.0	89.3	90.1	89.2	87.4	18.7
$K = 50$	91.5	89.7	90.9	89.2	90.3	88.5	89.3	86.9	88.7	10.3

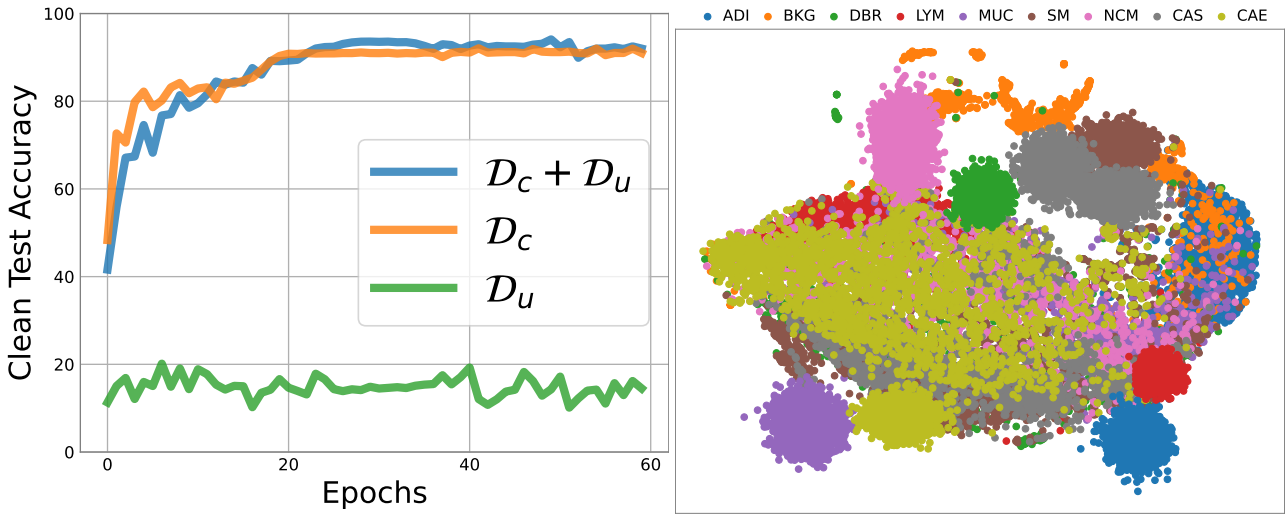

Figure 6. **Left:** The learning curve trained on RN-18 on three different datasets. **Right:** The t-SNE results of the whole mixed data.

on the labels. This spurred our investigation into the effectiveness of randomly selecting a subset of samples for perturbation. Specifically, we applied SALM to a selected percentage of the training data, leaving the remainder untouched and clean. We trained models on this mixed dataset of unlearnable and clean training data, $\mathcal{D}_u + \mathcal{D}_c$. For comparison, models are also trained on a completely clean dataset, denoted as $\mathcal{D}_c$.

The results in Table 7 show a rapid decrease in effectiveness when less than 100% of the data is selected for SALM application, whatever the K is. Surprisingly, applying SALM to as much as 80% of the data results in a negligible effect. Previous studies have shown that both error-minimizing noise and error-maximizing noise have similar limitations in DNNs (Huang et al., 2021; Shan et al., 2020).

To further illustrate this phenomenon, we take 80% as an example and plot the learning curves of ResNet-18 in the following scenarios: 1) Models trained on 20% clean data; 2) Models trained on 80% data processed by SALM; 3) Models trained on a mixture of these two types of data. The results are shown in Figure 6 left. It is observed that data processed by SALM at 80% remains unlearnable for the model, yet the mere 20% of clean data enables the model to achieve excellent performance. The superior performance trained on mixed data arises solely from the clean data contained within. Furthermore, the visualization results in Figure 6 reveal that the clean data remains mixed in 2D space, whereas the data processed by SALM is almost linearly separable, indicating its unlearnability. Therefore, in the case of sample-wise combined datasets, our method also demonstrates outstanding performance and robust scalability. Overall, our method provides a way for every collaborating entity within large datasets to protect their interests and ensure each data donor's contribution will not be misused.

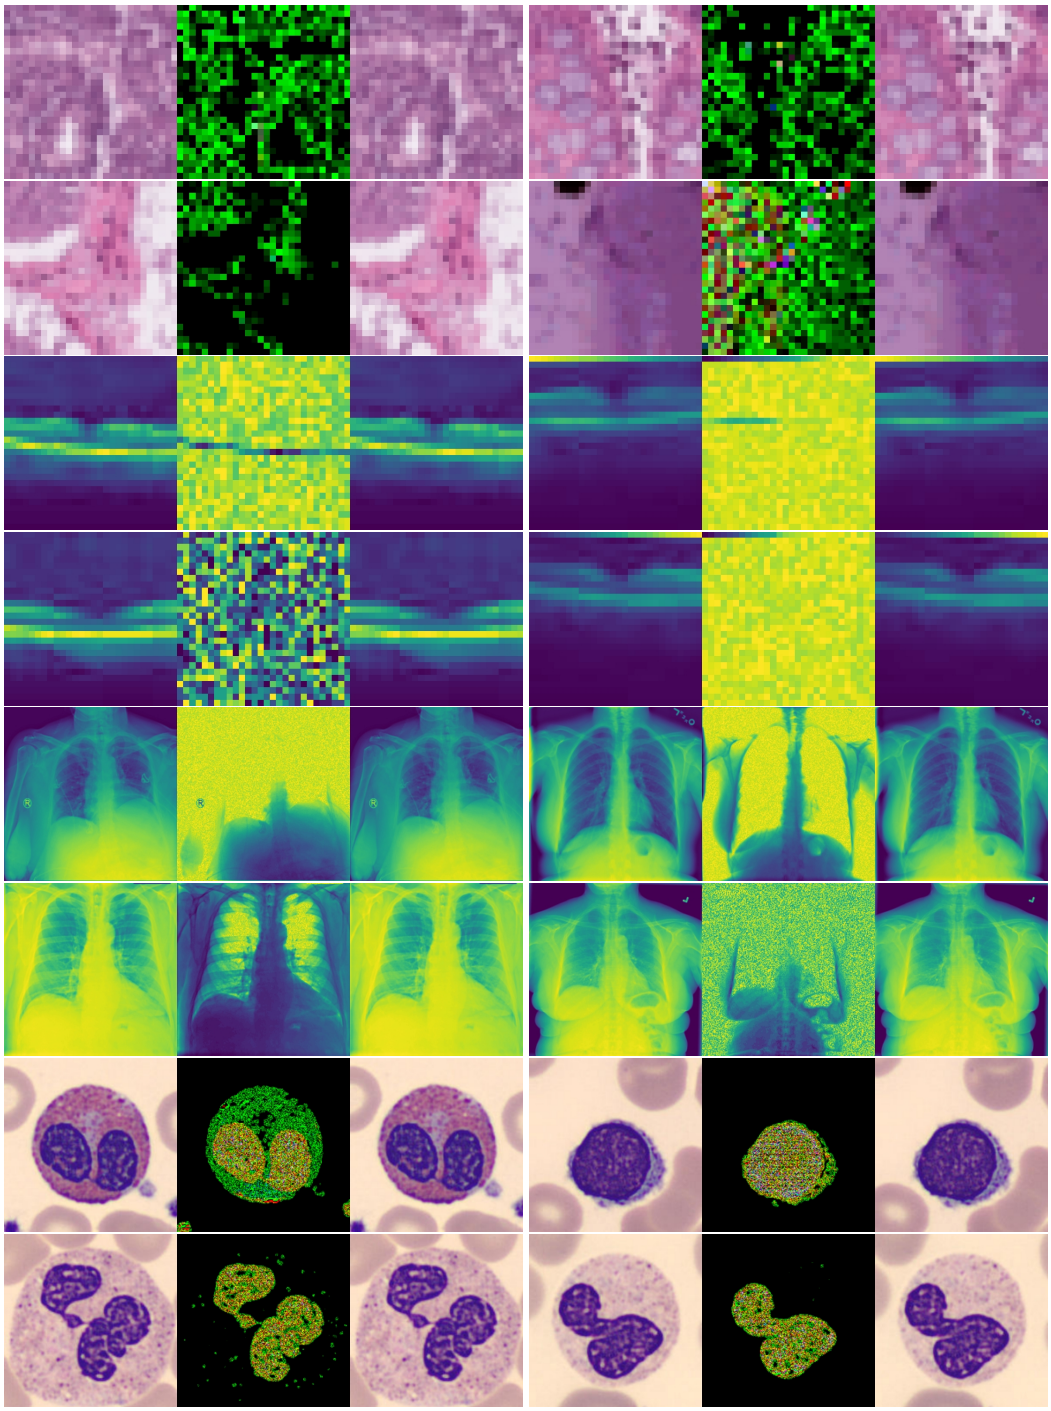

Figure 7. Some visualization results of the origin image and the corresponding noise and protected image.