

# 000 001 002 003 004 005 AUGMENTED MIXUP PROCEDURE FOR 006 PRIVACY-PRESERVING COLLABORATIVE TRAINING 007 008 009

010 **Anonymous authors**  
011 Paper under double-blind review  
012  
013  
014  
015  
016  
017  
018  
019  
020  
021  
022  
023  
024  
025  
026

## ABSTRACT

027 Mixup, introduced by Zhang et al., is a regularization technique for training neu-  
028 ral networks that generates convex combinations of input samples and their cor-  
029 responding labels. Motivated by this approach, Huang et al. proposed InstaHide,  
030 an image encryption method designed to preserve the discriminative properties  
031 of data while protecting original information during collaborative training across  
032 multiple parties. However, recent studies by Carlini et al., Luo et al., and Chen  
033 et al. have demonstrated that attacks exploiting the linear system generated by  
034 the mixup procedure can compromise the security guarantees of InstaHide. To  
035 address this vulnerability, we propose a modified mixing procedure that intro-  
036 duces perturbations into samples before forming convex combinations, making  
037 the associated linear inverse problem ill-conditioned for adversaries. We present a  
038 theoretical worst-case security analysis and empirically evaluate the performance  
039 of our method in mitigating such attacks. Our results indicate that robust attack  
040 mitigation can be achieved by increasing the perturbation level, without causing a  
041 significant reduction in classification accuracy. Furthermore, we compare the per-  
042 formance of our approach with that of InstaHide on standard benchmark datasets,  
043 including MNIST, CIFAR-10, CIFAR-100, and Tiny-ImageNet.  
044

## 1 INTRODUCTION

045 Data mixing was initially introduced as a dataset augmentation technique, generating new samples  
046 by computing weighted averages of subsets from the original dataset Zhang et al. (2017). Orig-  
047 inally designed as a regularization method for training neural networks, this approach has also been  
048 adapted for privacy-preserving protocols, as the mixing process obscures the original data during  
049 model training Liu et al. (2019); Fu et al. (2019).

050 Although the mixup strategy appears to preserve privacy without significantly degrading model per-  
051 formance, directly applying the method proposed by Zhang et al. (2017) can introduce vulnerabili-  
052 ties that allow attackers to recover private data under certain conditions Huang et al. (2020):  
053

- 054 1. **Mixup samples from a private dataset only:** If mixup samples are generated exclusively  
055 from a private dataset, an attacker can identify which samples share a common private com-  
056 ponent by analyzing the expected value of the dot product between mixup samples. Once a  
057 set of related mixup samples is identified, the common private sample can be reconstructed  
058 by averaging these samples.
- 059 2. **Mixup samples from both private and public datasets:** When mixup samples are gen-  
060 erated using both private and public datasets, repetitions of private samples in the mixup  
061 process can be avoided by leveraging the public dataset. However, since the public dataset  
062 is accessible, an attacker can perform a similar statistical analysis to identify which public  
063 samples were used in the mixup. Once the public components are determined, the private  
064 sample can be trivially reconstructed.

065 Building on their security analysis, Huang et al. (2020) proposed a mixup-based algorithm called  
066 InstaHide. The key innovation of their method is the application of a sign-flipping mask to images  
067 generated by computing a weighted sum of both public and private samples. The authors analyzed  
068 the security of InstaHide and formally proved that its security depends on the computational hard-  
069 ness of the subset-sum problem. However, the assumptions underlying their security model do not  
070

054 accurately reflect the properties of real-world data. In particular, Huang et al. (2020) assumed that  
 055 each sample consists of an arbitrary sequence of values. For example, in the context of images, this  
 056 assumption implies that pixels are independently and randomly distributed, which does not hold in  
 057 practice. Consequently, although the security proof is mathematically valid, it does not offer prac-  
 058 tical security guarantees. This limitation was highlighted by Carlini et al. (2021a), who developed  
 059 efficient attacks on samples generated by the InstaHide algorithm, enabling near-complete recovery  
 060 of the original data.

## 062 2 CONTRIBUTIONS

064 The primary contribution of this paper, presented in Section 3, is a singularized mixup algorithm that  
 065 resolves the core weakness exploited by major attacks on InstaHide: the repeated reuse of the same  
 066 private sample across mixtures. Our method mixes only two private images at a time and injects  
 067 structured noise into the non-target component, thereby eliminating the persistent signal required  
 068 for reconstruction attacks.

069 In Section 4, we develop a theoretical security analysis of the proposed mechanism. We establish  
 070 lower bounds on the achievable reconstruction error under an adversary with full knowledge of  
 071 the mixing weights, and we provide a principled way to choose the noise norm so that the SNR  
 072 associated with separating each encoded sample into signal and interference components remains  
 073 below a prescribed security parameter  $\tau$ .

074 Section 5 presents an empirical study of both security and accuracy. We evaluate linear and nonlinear  
 075 attackers in a conservative threat model and show that, above a modest noise threshold, neither can  
 076 recover the underlying private images. At the same time, even when using a noise level far stricter  
 077 than what the theoretical bounds require, our method maintains accuracy comparable to InstaHide’s  
 078 strongest  $k = 4$  configuration on MNIST, CIFAR-10, CIFAR-100, and Tiny-ImageNet. Overall, our  
 079 results demonstrate that the proposed mixup procedure provides strong protection against inversion  
 080 while preserving high downstream utility.

## 082 3 SINGULARIZED MIXUP

085 In this section, we introduce our algorithm, which is based on the singularization framework. This  
 086 approach enhances security by ensuring that each execution is unique, thereby making it broadly  
 087 applicable to a variety of systems Gaber et al. (2023). Singularization has previously been used to  
 088 strengthen encryption algorithms without modifying their underlying structure Macario-Rat & Plesa  
 089 (2024), which motivates our adoption of this framework in the design of our mixup algorithm.

090 We begin with a brief overview of InstaHide, followed by a detailed description of our proposed  
 091 algorithm.

### 093 3.1 INSTAHIDE ALGORITHM

095 Consider a private dataset  $(x_i, y_i)_{i=1}^n$  consisting of  $n$  samples, where  $x_i \in \mathbb{R}^d$  denotes the input  
 096 example and  $y_i \in \mathbb{R}^c$  is the corresponding one-hot encoded label.

097 The fundamental idea behind Mixup, as introduced by Zhang et al. (2017), is to replace each data  
 098 point with a convex combination of the current sample and  $k - 1$  other samples selected uniformly  
 099 at random from the dataset. Specifically, each new data point is generated by taking a weighted  
 100 average of  $k$  instances and their associated labels:

$$102 \quad \tilde{x}_i \leftarrow w_{i1}x_i + \sum_{j=2}^{k-1} w_{ij}x_{\pi_i(j)} \quad (1)$$

$$106 \quad \tilde{y}_i \leftarrow w_{i1}y_i + \sum_{j=2}^{k-1} w_{ij}y_{\pi_i(j)} \quad (2)$$

108 where  $\{(\tilde{x}_i, \tilde{y}_i)\}_{i=1}^n$  represents the encoded dataset and  $\pi_i$  is a random permutation over  
 109  $\{1, 2, \dots, n\}$ .  
 110

111 The InstaHide approach builds upon Mixup but introduces two key modifications:

112

- 113 1. **Public images:** InstaHide augments the private dataset with samples from public datasets,  
 114 expanding the pool of mixing samples to  $(x_i, y_i)_{i=1}^n \cup (x_i, y_i)_{i=n+1}^{n+m}$ , where  $m$  denotes the  
 115 size of the public dataset.
- 116 2. **Sign mask:** The sign of each pixel in a mixup image is randomly flipped using a random  
 117 sign mask  $\sigma_i \sim \Lambda_{\pm}^d$ .

118 As a result, equations (1) and (2) are modified as follows:

119

$$120 \tilde{x}_i \leftarrow \sigma_i \circ \left( w_{i1}x_i + \sum_{j=2}^{k_s-1} w_{ij}x_{\pi_i(j)} + \sum_{j=k_s+1}^{k-k_t} w_{ij}x_{\pi_{i_p}(j)} \right) \quad (3)$$

121

122

$$123 \tilde{y}_i \leftarrow w_{i1}y_i + \sum_{j=2}^{k_s-1} w_{ij}y_{\pi_i(j)} \quad (4)$$

124

125 Here,  $k_s$  denotes the number of private images,  $k_t$  the number of public images, and  $\pi_{i_p}$  is a random  
 126 permutation over the set  $\{n+1, \dots, n+m\}$ . Note that public images are used solely as a source of  
 127 structured noise, and their labels are not included in the mix.

### 132 3.2 SINGULARIZATION ALGORITHM

133 The principal aim of the singularization algorithm is to transform the original dataset  $\{(x_i, y_i)\}_{i=1}^n$   
 134 into a new set  $\{(\tilde{x}_i, \tilde{y}_i)\}_{i=1}^n$  such that the resulting dataset preserves the discriminative characteristics  
 135 of the original data, while ensuring that the original data cannot be recovered.

136 A key vulnerability of the InstaHide algorithm arises from the possibility that two encoded samples  
 137 may share the same private input during the mixup process. This stems from InstaHide's encoding  
 138 strategy, where each encoded sample is formed as a convex combination of two data points from the  
 139 original private dataset and  $k-2$  additional samples drawn from public sources. The inclusion of  
 140 multiple private samples in each encoded combination allows an attacker to group encoded samples  
 141 that share a common private component, making it possible to reconstruct the original private image  
 142 from these clusters. For the attacker, the repeated appearance of a private sample across different  
 143 encodings acts as a persistent signal, while the public components serve as noise that can be filtered  
 144 out through aggregation.

145 There are three primary distinctions between our approach and InstaHide. First, our method con-  
 146 structs each encoded input using exactly two private data points ( $k=2$ ), and does not incorporate  
 147 any public data. This design choice is motivated by security considerations. While InstaHide sug-  
 148 gests increasing  $k$  to mitigate brute-force attacks on the subset sum problem, Carlini et al. (2021a)  
 149 has demonstrated that larger values of  $k$  can actually reduce security. Specifically, when the mixing  
 150 weights are known, the resulting linear system becomes easier to solve, making it more vulne-  
 151 rable to attack. By fixing  $k=2$  and relying exclusively on private data, our approach avoids these  
 152 vulnerabilities.

153 Second, we introduce noise only to the second private data point before performing the mixup op-  
 154 eration. This strategy is intended to tightly couple the noise with the private information, thereby  
 155 making it significantly more difficult for an adversary to recover the original data. When the noise  
 156 level is sufficiently high, inverting the process becomes an ill-conditioned problem, which further  
 157 enhances security.

158 Finally, our method does not require sign-flipping masks or the use of public images. Previous  
 159 attacks Carlini et al. (2021a); Chen et al. (2020); Luo et al. (2022) have shown that sign-flipping  
 160 masks can be circumvented by analyzing the absolute values of mixed images, thus limiting their  
 161 effectiveness as a privacy mechanism. Our singularization algorithm is presented in Algorithm 1.

---

162 **Algorithm 1** Singularized Mixup

---

163 **Require:** Dataset  $\{(x_i, y_i)\}_{i=1}^n$ ; error norm  $r$

164 **Ensure:** Mixed dataset  $\{(\tilde{x}_i, \tilde{y}_i)\}_{i=1}^n$

165 1:  $\pi \sim \text{Uniform}(S_n)$

166 2: **for** each  $i = 1$  to  $n$  **do**

167 3:  $w_i \sim \text{Uniform}([0, 1]^2)$  and normalize such that  $\|w_i\|_1 = 1$  and  $\|w_i\|_\infty \leq \alpha$

168 4:  $e_i \sim \text{Uniform}(\mathbb{S}(0, r))$

169 5:  $\tilde{x}_i \leftarrow w_{i1}x_i + w_{i2}(x_{\pi(i)} + e_i)$

170 6:  $\tilde{y}_i \leftarrow w_{i1}y_i + w_{i2}y_{\pi(i)}$

171 7: **end for**

172 8: **return**  $\{(\tilde{x}_i, \tilde{y}_i)\}_{i=1}^n$

---

173

174

175 3.3 PRACTICAL INSTANTIATION

176 Similar to InstaHide Huang et al. (2020), the primary application of our algorithm is in privacy-  
 177 preserving collaborative training. Suppose there are multiple parties, each possessing a private local  
 178 dataset. These parties aim to jointly train a deep neural network on the combined data without  
 179 exposing the sensitive information contained in their individual datasets. The following general  
 180 framework demonstrates how Algorithm 1 can be integrated to achieve this goal:

181

- 182 1. All parties agree on a common preprocessing technique to be applied locally. For instance,  
 183 in the context of image data, participants may choose to standardize the images or extract  
 184 feature representations using a publicly available pretrained model, such as ResNet He et al.  
 185 (2016).
- 186 2. Each party independently transforms its local dataset by applying Algorithm 1, thereby  
 187 generating a set of mixup samples. Each sample consists of a mixup example and its  
 188 corresponding mixup label.
- 189 3. The resulting data is then transmitted to a central server, which is responsible for training  
 190 the deep learning model. Upon completion, the trained model is distributed back to the  
 191 parties for local use.

192

193 The security of this protocol depends on the effectiveness of Algorithm 1 in protecting the privacy  
 194 of local datasets. In particular, the central server, which only receives the mixup samples, should not  
 195 be able to reconstruct the original data from these representations.

196

197 4 SECURITY ANALYSIS

198 In this section, we analyze the security of the proposed scheme and evaluate its resilience to the  
 199 three principal attacks on the InstaHide framework, as presented in Carlini et al. (2021a), Chen et al.  
 200 (2020), and Luo et al. (2022). We begin by reviewing the attack strategy of Carlini et al. (2021a),  
 201 which applies directly to InstaHide without additional assumptions and forms the foundation for  
 202 subsequent attacks.

203 The attack of Carlini et al. (2021a) aims to recover the noisy linear system of equations generated  
 204 by the InstaHide encoding procedure. The adversary’s first task is to identify which private images  
 205 contribute to each encoded sample. To do so, the attacker generates encoded images using publicly  
 206 available data and trains a neural network to predict whether two encoded images share a common  
 207 private source image. Although InstaHide introduces random sign flips that might be expected  
 208 to impede this process, the authors show that taking the absolute value of each pixel effectively  
 209 removes this obstacle. Once trained, the comparison network allows the adversary to infer the  
 210 co-occurrence of private images across encoded samples. Since the mixing weights are revealed  
 211 through the encoded labels, the attacker can then assemble a noisy linear system whose noise arises  
 212 from the public images included in each mixup. Because the public images vary across samples,  
 213 this noise behaves approximately as mean-zero Gaussian and averages out across many equations.  
 214 Solving this system enables the attacker to reconstruct the private images up to a global sign per  
 215 pixel. A final recoloring step is applied to improve visual fidelity.

216 In contrast, the attack of Chen et al. (2020) relies on an explicit distributional assumption: the original images are modeled as Gaussian. Under this assumption, the absolute values of the encoded samples follow a folded Gaussian distribution. Given sufficiently many mixup samples, the adversary can estimate their covariance matrix, which corresponds to the Gram matrix of the mixing weight vectors. From this Gram matrix, the adversary can determine which private images participate in each mixup. As in Carlini et al. (2021a), this identification step enables the construction of a linear inverse problem whose solution yields the private images up to pixel-wise sign ambiguities.  
 217  
 218  
 219  
 220  
 221  
 222

223 To mitigate the attack of Carlini et al. (2021a), Luo et al. (2022) proposes introducing geometric  
 224 augmentations—such as shifting, cropping, rotation, and translation—to disrupt pixel-wise alignment  
 225 before mixup. This defense seeks to prevent the formation of a consistent linear system that  
 226 attackers could invert. Their approach shares a broad motivation with ours: both aim to ensure  
 227 that identical images do not reappear across multiple encoded samples. However, Luo et al. (2022)  
 228 demonstrates that such augmentations can be circumvented. Specifically, the attacker trains a  
 229 comparison network, as in Carlini et al. (2021a), to detect whether two encoded samples share a common  
 230 private image, even under geometric transformations. Encoded samples containing the same private  
 231 image are then clustered. Within each cluster, a fusion-denoising pipeline is applied: a convolutional  
 232 network first downsamples the encoded images to reduce geometric variability, a transpose  
 233 CNN upsamples the result, and multiple outputs are fused (by averaging or max-pooling) before  
 234 being passed through a denoising network that reconstructs the private image.  
 235

236 The attacks of Carlini et al. (2021a) and Chen et al. (2020) are ineffective against our method because  
 237 the injected noise disrupts the formation of stable linear systems. Moreover, unlike Luo et al. (2022),  
 238 our scheme does not rely on geometric transformations, thereby limiting the applicability of fusion-  
 239 denoising attacks. Although our mechanism is designed to prevent adversaries from forming clusters  
 240 or reconstructing underlying linear systems, we adopt a conservative security model in which the  
 241 adversary is assumed capable of doing so. This assumption is motivated by a common feature of  
 242 all three attacks: the ability to determine whether two encoded images share a private component.  
 243 In both Carlini et al. (2021a) and Chen et al. (2020), the random sign mask is rendered ineffective  
 244 by taking absolute values, while in Luo et al. (2022), geometric transformations do not prevent  
 245 clustering.

246 To assess the security of our singularized mixup algorithm, we characterize the reconstruction error  
 247 faced by an adversary attempting to invert the system of equations induced by Algorithm 1. We  
 248 assume the attacker has access to both the encoded samples and the mixing weights, and seeks to  
 249 recover the original data. In Theorem 4.1, we examine an adversary that does not leverage any struc-  
 250 tural prior, such as assuming the unknowns are images or adhere to particular statistical constraints.  
 251 The theorem shows that the expected Euclidean recovery error scales linearly with the noise radius  
 252  $r$ .  
 253

254 **Theorem 4.1.** *Let  $X \in \mathbb{R}^{n \times d}$  have rows  $x_i^\top$ . Algorithm 1 produces*

$$\tilde{x}_i = w_{i1} x_i + w_{i2} (x_{\pi(i)} + e_i), \quad e_i \stackrel{i.i.d.}{\sim} \text{Uniform}(\mathbb{S}(0, r)),$$

255 *with  $\|w_i\|_1 = 1$  and  $\|w_i\|_\infty \leq \alpha$ . Let  $P_\pi$  be the permutation matrix of  $\pi$  and define*

$$256 \quad D_1 := \text{diag}(w_{11}, \dots, w_{n1}), \quad D_2 := \text{diag}(w_{12}, \dots, w_{n2}), \quad W := D_1 + D_2 P_\pi,$$

257 *and assume  $W$  is invertible. Then*

$$258 \quad \tilde{X} = W X + E, \quad E_i = w_{i2} e_i.$$

259 *For any estimator  $\hat{x}_i = \hat{x}_i(\tilde{X}, W)$ ,*

$$260 \quad \sup_{X \in \mathbb{R}^{n \times d}} \mathbb{E}[\|x_i - \hat{x}_i(\tilde{X}, W)\|_2^2] \geq r^2 T_i,$$

261 *where*

$$262 \quad T_i := \sum_{\ell=1}^n (W^{-1})_{i\ell}^2 w_{\ell2}^2.$$

263 *Proof.* Since  $W$  is invertible, define

$$264 \quad Y := W^{-1} \tilde{X} = X + Z, \quad Z := W^{-1} E.$$

270 Because  $E_\ell = w_{\ell 2} e_\ell$  and  $e_\ell$  are independent, mean-zero, and isotropic with  
 271

$$272 \quad \text{Cov}(e_\ell) = \frac{r^2}{d} I_d,$$

274 we have

$$275 \quad \text{Cov}(E_\ell) = w_{\ell 2}^2 \frac{r^2}{d} I_d.$$

277 Thus

$$278 \quad \text{Cov}(Z_i) = \sum_{\ell=1}^n (W^{-1})_{i\ell}^2 \text{Cov}(E_\ell) = \frac{r^2}{d} T_i I_d,$$

280 and hence

$$281 \quad \mathbb{E}\|Z_i\|_2^2 = \text{tr}(\text{Cov}(Z_i)) = r^2 T_i.$$

283 Fix any estimator  $\hat{x}_i(\tilde{X}, W)$  and write  $\delta(Y) := \hat{x}_i(\tilde{X}, W)$ . Consider matrices  $X$  with all rows equal  
 284 to zero except the  $i$ th. For such  $X$ ,

$$286 \quad Y_i = x_i + Z_i, \quad (Y_{-i}, Z_i) \text{ independent of } x_i.$$

287 Let  $m(Y_i) := \mathbb{E}[\delta(Y) | Y_i]$ . By Jensen's inequality,

$$289 \quad \mathbb{E}[\|x_i - m(Y_i)\|_2^2] \leq \mathbb{E}[\|x_i - \delta(Y)\|_2^2] \quad \text{for every } x_i \in \mathbb{R}^d.$$

291 Now restrict attention to this one-dimensional family  $X$  parametrized by  $x_i$ . We obtain the  $d$ -  
 292 dimensional location model

$$293 \quad Y_i = x_i + Z_i, \quad x_i \in \mathbb{R}^d,$$

294 with fixed noise  $Z_i$ . A standard Bayesian lower bound for location models implies

$$295 \quad \sup_{x_i \in \mathbb{R}^d} \mathbb{E}[\|x_i - m(Y_i)\|_2^2] \geq \mathbb{E}\|Z_i\|_2^2.$$

298 Combining the inequalities,

$$299 \quad \sup_X \mathbb{E}[\|x_i - \hat{x}_i(\tilde{X}, W)\|_2^2] \geq \mathbb{E}\|Z_i\|_2^2 = r^2 T_i.$$

301 This completes the proof.  $\square$

303 While Theorem 4.1 provides a tight minimax lower bound on the MSE for estimators that lack prior  
 304 knowledge, it does not prescribe how to choose the noise norm  $r$ . Our goal is to select  $r$  such  
 305 that encoded samples do not leak meaningful information about the originals. We formalize this  
 306 through the signal-to-noise ratio (SNR) associated with the decomposition of each encoded sample  
 307 into signal and interference components. Motivated by the defense rationale in Luo et al. (2022),  
 308 where distinct augmented images are used to prevent redundancy in mixup inputs, we choose  $r$   
 309 proportional to the typical separation between data points. Theorem 4.2 specifies how to select the  
 310 scaling parameter  $mf$  and thereby set  $r$  as a multiple of the average inter-sample distance so that the  
 311 resulting SNR falls below a prescribed threshold  $\tau$ .

312 **Theorem 4.2.** *Let  $\{x_i\}_{i=1}^n \subset \mathbb{R}^d$  be i.i.d. samples of a subgaussian random vector  $X$  with  $\mathbb{E}[X] = 0$   
 313 and covariance  $\Sigma$ . Define*

$$314 \quad V = \mathbb{E}[\|X\|_2^2] = \text{tr}(\Sigma), \quad D = \mathbb{E}[\|X - X'\|_2],$$

316 for an independent copy  $X'$ . Set

$$317 \quad c := \frac{D^2}{2V} \in (0, 1],$$

319 so that  $D^2 = 2cV$  (Jensen and independence give  $D^2 \leq 2V$ ). Algorithm 1 outputs, for each  $i$ ,

$$321 \quad \tilde{x}_i = w_{i1} x_i + (1 - w_{i1})(x_{\pi(i)} + e_i),$$

322 where  $0 \leq w_{i1} \leq \alpha$  almost surely for a given  $\alpha \in (0, 1)$ , and the error vectors satisfy

$$323 \quad \|e_i\|_2 = r, \quad r = mf D.$$

324 Assume that  $w_{i1}$  is independent of the data and that  $e_i$  is independent of the data and of  $w_{i1}$ . Define  
 325 the signal and interference components

$$326 \quad S_i = w_{i1} x_i, \quad I_i = (1 - w_{i1})(x_{\pi(i)} + e_i),$$

327 and the expected signal-to-noise ratio

$$329 \quad \text{SNR} = \frac{\mathbb{E}\|S_i\|_2^2}{\mathbb{E}\|I_i\|_2^2}.$$

331 Then

$$332 \quad \text{SNR} \leq \frac{\alpha^2}{(1 - \alpha)^2 \left(1 + \frac{r^2}{V}\right)} = \frac{\alpha^2}{(1 - \alpha)^2 \left(1 + \frac{m_f^2 D^2}{V}\right)} = \frac{\alpha^2}{(1 - \alpha)^2 \left(1 + 2c m_f^2\right)}.$$

336 Consequently, for any target  $\tau > 0$ ,

$$337 \quad \text{SNR} \leq \tau \quad \text{whenever} \quad m_f \geq \sqrt{\frac{1}{2c} \left( \frac{\alpha^2}{\tau(1 - \alpha)^2} - 1 \right)}.$$

340 *Proof.* Since  $0 \leq w_{i1} \leq \alpha$  a.s. and  $w_{i1}$  is independent of  $x_i$ ,

$$342 \quad \mathbb{E}\|S_i\|_2^2 = \mathbb{E}[w_{i1}^2 \|x_i\|_2^2] = \mathbb{E}[w_{i1}^2] \mathbb{E}[\|X\|_2^2] \leq \alpha^2 V.$$

343 For the interference, using independence of  $w_{i1}, X', e_i$ , zero-mean  $\mathbb{E}[X'] = 0$ , and the polarization  
 344 identity,

$$345 \quad \mathbb{E}\|I_i\|_2^2 = \mathbb{E}[(1 - w_{i1})^2 \|X' + e_i\|_2^2] \\ 346 \quad = \mathbb{E}[(1 - w_{i1})^2] \mathbb{E}[\|X'\|_2^2] + \mathbb{E}[(1 - w_{i1})^2] \mathbb{E}[\|e_i\|_2^2] + 2\mathbb{E}[(1 - w_{i1})^2] \mathbb{E}[\langle X', e_i \rangle].$$

348 The cross term vanishes since  $\mathbb{E}[X'] = 0$  and  $e_i$  is independent of  $X'$ . Hence

$$349 \quad \mathbb{E}\|I_i\|_2^2 = \mathbb{E}[(1 - w_{i1})^2] (V + \mathbb{E}\|e_i\|_2^2).$$

350 Because  $0 \leq w_{i1} \leq \alpha$  a.s., we have  $(1 - w_{i1}) \geq (1 - \alpha)$  a.s., whence

$$352 \quad \mathbb{E}[(1 - w_{i1})^2] \geq (1 - \alpha)^2.$$

353 By construction  $\|e_i\|_2 = r$  deterministically, so  $\mathbb{E}\|e_i\|_2^2 = r^2$ . Therefore,

$$354 \quad \mathbb{E}\|I_i\|_2^2 \geq (1 - \alpha)^2 (V + r^2).$$

355 Taking the ratio gives

$$357 \quad \text{SNR} = \frac{\mathbb{E}\|S_i\|_2^2}{\mathbb{E}\|I_i\|_2^2} \leq \frac{\alpha^2 V}{(1 - \alpha)^2 (V + r^2)} = \frac{\alpha^2}{(1 - \alpha)^2 (1 + \frac{r^2}{V})}.$$

359 Substituting  $r = m_f D$  yields

$$361 \quad \text{SNR} \leq \frac{\alpha^2}{(1 - \alpha)^2 \left(1 + \frac{m_f^2 D^2}{V}\right)} = \frac{\alpha^2}{(1 - \alpha)^2 \left(1 + 2c m_f^2\right)},$$

363 because  $D^2 = 2cV$  by definition of  $c$ . Finally, solving

$$365 \quad \frac{\alpha^2}{(1 - \alpha)^2 \left(1 + 2c m_f^2\right)} \leq \tau$$

368 for  $m_f$  gives

$$369 \quad m_f \geq \sqrt{\frac{1}{2c} \left( \frac{\alpha^2}{\tau(1 - \alpha)^2} - 1 \right)},$$

371 which completes the proof.  $\square$

373 Finally, although Theorem 4.1 characterizes the fundamental difficulty of inversion for adversaries  
 374 without prior information, practical attackers may possess substantial knowledge about the structure  
 375 or distribution of the underlying data. To bridge this gap, our experimental evaluation assesses the  
 376 performance of both linear and nonlinear estimators in realistic scenarios, thereby providing a more  
 377 comprehensive understanding of the scheme's robustness against adversaries capable of exploiting  
 informative priors.

378 

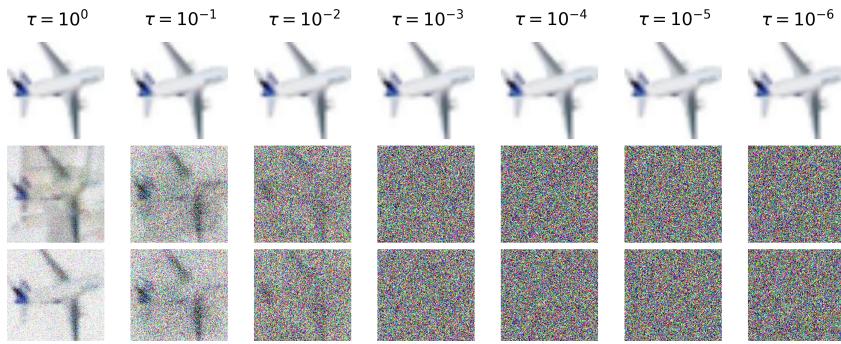
## 5 EXPERIMENTS

380 Our experiments are designed to evaluate both the accuracy loss relative to InstaHide and the attack  
381 resilience of our proposed algorithm.  
382383 **Setup:** We conduct our experiments on three widely used benchmark datasets: MNIST LeCun  
384 (1998), CIFAR-10, CIFAR-100 Krizhevsky et al. (2009) and Tiny-ImageNet Le & Yang (2015).  
385 All implementations are carried out using the PyTorch framework Paszke et al. (2019). For the  
386 classification accuracy experiments, we utilize feature representations obtained from the output of  
387 the final convolutional layer of publicly available pretrained image models. Specifically, feature  
388 maps are extracted using a pretrained ResNet-18 model for MNIST and CIFAR-10, and a pretrained  
389 ResNet-50 model for CIFAR-100 and Tiny-ImageNet.  
390391 **Security:**  
392393 In the following, we examine two distinct attack strategies for recovering the underlying images from  
394 their mixed representations. The first is a linear inversion attack, which exploits the known linear  
395 mixing process to directly reconstruct the sources. The second is a non-linear reconstruction attack  
396 based on U-Net architectures, allowing the adversary to learn a more flexible, data-driven inverse  
397 mapping. We evaluate the quality of the recovered images using the SNR for a range of values of  
398  $\tau$ , with the noise norm chosen according to Theorem 4.2. In particular, we estimate the average  
399 distance between two randomly selected images directly from the data and set the noise level  $r$  to  
400 be this empirical average multiplied by the scaling factor prescribed by the theorem. This procedure  
401 ensures that the noise magnitude is consistent with the theoretical regime under consideration.  
402403 For the linear inversion attack, we follow a gradient-descent-based reconstruction procedure in-  
404 spired by Luo et al. (2022). In this setting, the adversary is assumed to know the linear mixing  
405 matrix and uses this information to guide the recovery process. The loss function is built around a  
406 linear reconstruction term that measures the mean squared error between the observed mixtures and  
407 the linearly recomposed images obtained by applying the known mixing weights to the current esti-  
408 mates. To stabilize this inversion, the attacker incorporates generic priors on natural images through  
409 additional regularization terms: a total-variation penalty that promotes spatial smoothness and  $L_2$   
410 penalty that discourages unrealistically large pixel values. These regularizers encode broad assump-  
411 tions about natural images—namely smoothness and bounded intensity—without introducing any  
412 nonlinear modeling. Optimization is carried out with Adam, and after each update the recovered  
413 images are clipped to a fixed range to maintain plausible pixel values. The resulting reconstruction  
414 performance is reported in Table 1, and Figure 3 shows the best recovered image for CIFAR-10  
415 under the linear attack. Results for all datasets are provided in Appendix G.  
416417 

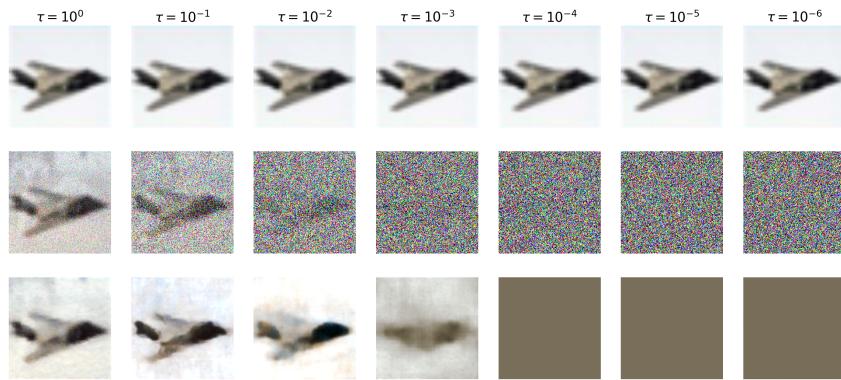
Table 1: Average reconstruction SNR (dB)  $\pm$  standard deviation for the linear attack for different  
418 noise levels  $\tau$ .

$\tau$	<b>MNIST</b>	<b>CIFAR-10</b>	<b>CIFAR-100</b>	<b>Tiny-ImageNet</b>
$10^0$	$8.19 \pm 0.31$	$0.85 \pm 2.69$	$0.94 \pm 3.01$	$1.24 \pm 2.47$
$10^{-1}$	$0.09 \pm 0.14$	$-5.06 \pm 2.29$	$-4.86 \pm 2.52$	$-4.60 \pm 2.01$
$10^{-2}$	$-2.79 \pm 0.09$	$-6.92 \pm 2.10$	$-6.66 \pm 2.30$	$-6.42 \pm 1.79$
$10^{-3}$	$-3.65 \pm 0.08$	$-7.45 \pm 2.05$	$-7.17 \pm 2.24$	$-6.94 \pm 1.75$
$10^{-4}$	$-3.90 \pm 0.09$	$-7.61 \pm 2.03$	$-7.32 \pm 2.22$	$-7.09 \pm 1.73$
$10^{-5}$	$-3.98 \pm 0.09$	$-7.66 \pm 2.03$	$-7.37 \pm 2.22$	$-7.14 \pm 1.73$
$10^{-6}$	$-4.01 \pm 0.09$	$-7.67 \pm 2.03$	$-7.38 \pm 2.22$	$-7.15 \pm 1.73$

423 For the nonlinear attack, we replace the analytic inversion used in the linear setting with a learned  
424 reconstruction model based on a U-Net architecture. The adversary is assumed to have access to  
425 a collection of public images, which are used to train the network to map mixed inputs back to  
426 clean sources. In our experiments, Tiny-ImageNet serves as the public dataset for training, while  
427 CIFAR-10 plays the role of the private dataset used for evaluation. A separate U-Net is trained  
428 for each value of  $\tau$ , ensuring that the attacker can adapt to the corresponding noise magnitude.  
429 Training is performed using an  $\ell_1$  reconstruction loss together with total-variation regularization,  
430 and predictions are kept within valid intensity ranges through standard normalization and clamping.  
431 The resulting reconstruction SNR for all  $\tau$  values are reported in Table 2. Figure 2 shows the best  
432 recovered CIFAR-10 example across noise levels.  
433

Figure 1: Recovered images under the linear attack for decreasing noise levels  $\tau$ Table 2: Average reconstruction SNR (dB)  $\pm$  standard deviation for the nonlinear attack across noise levels  $\tau$ .

$\tau$	$10^0$	$10^{-1}$	$10^{-2}$	$10^{-3}$	$10^{-4}$	$10^{-5}$	$10^{-6}$
	$7.76 \pm 3.46$	$7.27 \pm 3.39$	$5.44 \pm 2.93$	$2.84 \pm 2.28$	$-0.13 \pm 0.54$	$-0.13 \pm 0.53$	$-0.14 \pm 0.58$

Figure 2: Nonlinear (U-Net) attack on CIFAR-10: ground truth, mixed inputs, and recovered images across noise levels  $\tau$ **Accuracy:**

The security analysis indicates that, although the nonlinear attack is noticeably more powerful than the linear one for  $\tau \geq 10^{-4}$ , the original image remains unrecoverable across all datasets even at this relatively small noise level. Since  $\tau = 10^{-4}$  already suffices to prevent meaningful reconstruction, we adopt a conservative stance and conduct all accuracy experiments at an even stricter privacy setting, using  $\tau = 10^{-6}$  uniformly across datasets. At this noise level, we train a standard convolutional classifier on mixed representations and report the resulting test accuracy for each dataset. We then compare these results with the best-performing configuration of InstaHide (with  $k = 4$ ) to quantify the utility–privacy trade-off under a conservative noise regime. The full accuracy results are summarized in Table 3.

Table 3: Test accuracy (%) at  $\tau = 10^{-6}$  compared with the best reported InstaHide configuration ( $k = 4$ ).

Method	MNIST	CIFAR-10	CIFAR-100	Tiny-ImageNet
InstaHide ( $k = 4$ )	99.66	91.20	74.01	–
Ours ( $\tau = 10^{-6}$ )	99.32	90.51	75.99	72.50

The accuracy results in Table 3 show that, even under the conservative noise setting  $\tau = 10^{-6}$ —significantly stricter than what is needed to prevent both linear and nonlinear reconstruction—the loss in predictive performance remains negligible across all datasets. On MNIST and

486 CIFAR-10, our approach matches the accuracy of InstaHide with  $k = 4$ , and on CIFAR-100 and  
 487 Tiny-ImageNet it even yields modest improvements. These findings indicate that strong recon-  
 488 struction resistance does not come at the expense of meaningful degradation in downstream util-  
 489 ity: despite enforcing a noise level far below the threshold where the nonlinear attacker fails (i.e.,  
 490  $\tau \approx 10^{-4}$ ), the mixed representations still support high classification accuracy. Overall, the  
 491 results demonstrate that our method provides robust security guarantees while preserving competitive  
 492 model performance across diverse datasets.

493 We additionally evaluate our method in a federated learning setting by partitioning the CIFAR-10  
 494 dataset across multiple parties. Each party holds a disjoint local subset of the data. We compare (i) a  
 495 baseline model trained solely on the local subset of Party 0, and (ii) a model trained on the union of  
 496 all mixed datasets produced by Algorithm 1 across the participating parties. The results in Table 4  
 497 show that training on the aggregated mixed representations consistently achieves higher accuracy as  
 498 the number of parties increases.

499  
 500 Table 4: Federated CIFAR-10 test accuracy (%). The baseline model is trained only on Party 0’s  
 501 local data, while the second model is trained on the union of mixed datasets produced by all parties  
 502 using Algorithm 1.

503 <b>Number of Parties</b>	504 <b>Baseline (Local Only)</b>	505 <b>Mixup-Union (All Parties)</b>
504 10	505 87.25	506 90.63
505 20	506 85.56	507 90.32
506 30	507 84.08	508 88.23

## 509 6 CONCLUSIONS

510 In this paper, we introduced a singularized mixup mechanism that mixes only two private images at  
 511 a time while injecting structured noise into the non-target component. This design directly addresses  
 512 the key vulnerability exploited in attacks on InstaHide, where repeated use of the same private image  
 513 across many mixtures enables clustering and subsequent reconstruction. By corrupting all but one  
 514 component in each mixture, our method prevents such clustering and limits what an adversary can  
 515 infer from any single mixed sample.

516 We analyzed security under a conservative threat model in which the attacker has full knowledge  
 517 of the mixing weights and evaluated both a linear inversion attack and a more powerful nonlinear  
 518 U-Net-based attack. Our experiments show that once the noise magnitude exceeds a modest thresh-  
 519 old, neither attack can recover the original image, even when given maximal information about the  
 520 mixing process. At the same time, using a conservative setting of  $\tau = 10^{-6}$ —well below the level  
 521 at which both attacks already fail—the mixed representations retain high utility, with only negligible  
 522 accuracy loss compared to InstaHide’s best  $k = 4$  configuration.

523 While the method is tailored to image data, the singularization principle may extend to other modal-  
 524 ities with appropriate noise models. Exploring such extensions presents an interesting direction for  
 525 future work toward broadly applicable, attack-resistant data mixing schemes.

## 527 528 REPRODUCIBILITY STATEMENT

529 All source code used in this study is publicly available, along with detailed instructions to reproduce  
 530 the experiments described in this paper. The data utilized comes from publicly accessible benchmark  
 531 datasets. For additional details regarding the experimental setup and procedures, please refer to the  
 532 Appendix E.

533  
 534  
 535  
 536  
 537  
 538  
 539

540 REFERENCES  
541

542 Nicholas Carlini, Samuel Deng, Sanjam Garg, Somesh Jha, Saeed Mahloujifar, Mohammad Mah-  
543 moody, Abhradeep Thakurta, and Florian Tramèr. Is private learning possible with instance en-  
544 coding? In *2021 IEEE Symposium on Security and Privacy (SP)*, pp. 410–427. IEEE, 2021a.

545 Nicholas Carlini, Sanjam Garg, Somesh Jha, Saeed Mahloujifar, Mohammad Mahmoody, and Flo-  
546 rian Tramer. Neuracrypt is not private. *arXiv preprint arXiv:2108.07256*, 2021b.

547 Sitan Chen, Xiaoxiao Li, Zhao Song, and Danyang Zhuo. On instahide, phase retrieval, and sparse  
548 matrix factorization. *arXiv preprint arXiv:2011.11181*, 2020.

549 Shaltiel Eloul, Fran Silavong, Sanket Kamthe, Antonios Georgiadis, and Sean J Moran. Mixing  
550 gradients in neural networks as a strategy to enhance privacy in federated learning. In *Proceedings  
551 of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 3956–3965, 2024.

552 Yingwei Fu, Huaimin Wang, Kele Xu, Haibo Mi, and Yijie Wang. Mixup based privacy preserving  
553 mixed collaboration learning. In *2019 IEEE International Conference on Service-Oriented System  
554 Engineering (SOSE)*, pp. 275–2755. IEEE, 2019.

555 Chrystel Gaber, Gilles Macariot-Rat, Simona David, Jean-Philippe Wary, and Alain Cuaboz. Posi-  
556 tion paper: Strengthening applets on legacy sim cards with singularization, a new moving target  
557 defense strategy. In *International Conference on Mobile, Secure, and Programmable Networking*,  
558 pp. 71–74. Springer, 2023.

559 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recog-  
560 nition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition  
(CVPR)*, pp. 770–778, 2016.

561 Qijian He, Wei Yang, Bingren Chen, Yangyang Geng, and Liusheng Huang. Transnet: Training  
562 privacy-preserving neural network over transformed layer. *Proceedings of the VLDB Endowment*,  
563 13(12):1849–1862, 2020.

564 Yangsibo Huang, Zhao Song, Kai Li, and Sanjeev Arora. Instahide: Instance-hiding schemes for  
565 private distributed learning. In *International conference on machine learning*, pp. 4507–4518.  
566 PMLR, 2020.

567 Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny im-  
568 ages.(2009), 2009.

569 Yann Le and Xuan Yang. Tiny imagenet visual recognition challenge. *CS 231N*, 7(7):3, 2015.

570 Yann LeCun. The mnist database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>, 1998.

571 Q Li, Y Zhang, J Ren, Q Li, and Y Zhang. You can use but cannot recognize: Preserving visual  
572 privacy in deep neural networks. arxiv 2024. *arXiv preprint arXiv:2404.04098*.

573 Zhijian Liu, Zhanghao Wu, Ligeng Zhu, Chuang Gan, and Song Han. Facemix: Privacy-preserving  
574 face attribute classification on the cloud. 2019.

575 Zhijian Liu, Zhanghao Wu, Chuang Gan, Ligeng Zhu, and Song Han. Datamix: Efficient privacy-  
576 preserving edge-cloud inference. In *Computer Vision–ECCV 2020: 16th European Conference,  
577 Glasgow, UK, August 23–28, 2020, Proceedings, Part XI 16*, pp. 578–595. Springer, 2020.

578 Xinjian Luo, Xiaokui Xiao, Yuncheng Wu, Juncheng Liu, and Beng Chin Ooi. A fusion-denoising  
579 attack on instahide with data augmentation. In *Proceedings of the AAAI Conference on Artificial  
580 Intelligence*, volume 36, pp. 1899–1907, 2022.

581 Gilles Macario-Rat and Mihail-Iulian Plesa. Singularization: A new approach to designing block  
582 ciphers for resource-constrained devices. In *International Conference on Attacks and Defenses  
583 for Internet-of-Things*, pp. 155–167. Springer, 2024.

584 S Nythia, S Priyaa, R Priyanka, and S Saranya. Face image recognition and scrambling for privacy  
585 using neural networks. *International Journal of Scientific & Engineering Research (IJE), India*,  
586 2017.

594 Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor  
 595 Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward  
 596 Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner,  
 597 Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance  
 598 deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché Buc, E. Fox,  
 599 and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 32, pp. 8024–  
 600 8035. Curran Associates, Inc., 2019.

601 Andreea Bianca Popescu, Ioana Antonia Taca, Anamaria Vizitiu, Cosmin Ioan Nita, Constantin  
 602 Suciu, Lucian Mihai Itu, and Alexandru Scafa-Udriste. Obfuscation algorithm for privacy-  
 603 preserving deep learning-based medical image analysis. *Applied Sciences*, 12(8):3997, 2022.

604 Sifat Ut Taki and Spyridon Mastorakis. Amalgam: A framework for obfuscated neural network  
 605 training on the cloud. In *Proceedings of the 25th International Middleware Conference*, pp. 238–  
 606 251, 2024.

607 Yinggui Wang, Yuanqing Huang, Jianshu Li, Le Yang, Kai Song, and Lei Wang. Adaptive hybrid  
 608 masking strategy for privacy-preserving face recognition against model inversion attack. *arXiv*  
 609 preprint *arXiv:2403.10558*, 2024.

610 Yuxin Xiang, Tiantian Li, Wei Ren, Tianqing Zhu, and Kim-Kwang Raymond Choo. A lightweight  
 611 privacy-preserving scheme using pixel block mixing for facial image classification in deep learning.  
 612 *Engineering Applications of Artificial Intelligence*, 126:107180, 2023.

613 Hanshen Xiao, G Edward Suh, and Srinivas Devadas. Formal privacy proof of data encoding: The  
 614 possibility and impossibility of learnable encryption. In *Proceedings of the 2024 on ACM SIGSAC*  
 615 Conference on Computer and Communications Security

616 Conference on Computer and Communications Security

617 Conference on Computer and Communications Security

618 A Yala et al. Neuracrypt: hiding private health data via random neural networks for public training  
 619 (2021).

620 Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical  
 621 risk minimization. *arXiv preprint arXiv:1710.09412*, 2017.

## 624 A CHEN ATTACK

625 The attack described in Chen et al. (2020) simplifies the InstaHide problem by assuming that the  
 626 matrix  $X \in \mathbb{R}^{d \times n}$  of images is Gaussian, i.e., its entries are chosen i.i.d. from  $\mathcal{N}(0, 1)$ . Let  
 627  $p_1, \dots, p_d \in \mathbb{R}^n$  be the rows of  $X$ . Consider  $w_{i1}, \dots, w_m \in \mathbb{R}^n$  the unknown selection vectors  
 628 chosen from a distribution  $\mathcal{D}$ .  $S \subset \{1, \dots, m\}$  be the coordinates of the public images and  $S^c =$   
 629  $\{1, \dots, n\} \setminus S$  be the coordinates of the private images. Let  $[v]_S \in \mathbb{R}^{|S|}$  be the restriction of a vector  
 630  $v$  to the coordinates indexed by  $S$ . Each selection sector generates an encoded image as:

$$633 \quad 634 \quad \tilde{x}_i = |Xw_i| \quad (5)$$

635 In InstaHide, the sign of each pixel from an encoded image is randomly flipped, but as the authors  
 636 remark, the two notations are interchangeable.

637 The attack goes like follows:

639 1. **Learning the public coordinates of any selection vector** In the first step, the attacker  
 640 determines the weights associated with the public images from each selection vector. Con-  
 641 sidering the matrix

$$642 \quad 643 \quad N = \mathbb{E}_{p, \tilde{x}} [\tilde{x}^2 \cdot ([p]_S [p]_S^\top - \text{Id})]$$

644 where  $\tilde{x} = |\langle w, p \rangle|$ ,  $p \sim \mathcal{N}(0, \text{Id})$ . It can be proven that  $N$  is a rank-1 matrix proportional  
 645 to  $[w]_S [w]_S^\top$ . Moreover,  $N$  can be approximated by

$$646 \quad 647 \quad \hat{N} = \frac{1}{d} \sum_{i=1}^d \tilde{x}_i^2 \cdot ([p_i]_S [p_i]_S^\top - \text{Id}).$$

648  
 649  
 650  
 651  
 652  
 653  
 654  
 655  
 656  
 657  
 658  
 659  
 660  
 661  
 662  
 663  
 664  
 665  
 666  
 667  
 668  
 669  
 670  
 671  
 672  
 673  
 674  
 675  
 676  
 677  
 678  
 679  
 680  
 681  
 682  
 683  
 684  
 685  
 686  
 687  
 688  
 689  
 690  
 691  
 692  
 693  
 694  
 695  
 696  
 697  
 698  
 699  
 700  
 701

2. **Recovering the Gram Matrix** Since the previous steps recover the coordinates of the public images in each selection vector, for the simplicity of the description, consider that all images are private, i.e.,  $S^c = \{1, \dots, n\}$ . Consider the matrix  $\tilde{X} \in \mathbb{R}^{m \times d}$  where each line is an encoded images:

$$\tilde{X} = \begin{pmatrix} |\langle p_1, w_1 \rangle| & \cdots & |\langle p_d, w_1 \rangle| \\ \vdots & \ddots & \vdots \\ |\langle p_1, w_m \rangle| & \cdots & |\langle p_d, w_m \rangle| \end{pmatrix} \quad (6)$$

We can use the columns of  $\tilde{X}$  to estimate the covariance matrix  $\tilde{M}$  of the folded Gaussian distribution  $\mathcal{N}^{\text{fold}}(0, M)$ , since each column is drawn independently from this distribution. The covariance matrix  $M$  is, in fact, the rescaled  $m \times m$  Gram matrix whose entries are proportional to the dot product of any two selection vectors; that is, the element at position  $(i, j)$  in the matrix  $M$  is given by  $k \cdot \langle w_i, w_j \rangle$ .

### 3. Floral submatrices

The previous step of the attack shows the dot product between any two selection vectors, i.e.,  $\langle w_i, w_j \rangle$ , thus the attacker knows how many private images are common between two encoded images. In order to identify which private images are common (not only how many), the attacker identifies  $M$  floral submatrices. The rows/columns of a floral submatrix can be indexed by all subsets of size  $k$  of a set of  $k+2$  elements where its entries are the intersection sizes between the subsets. More intuitively, the attacker exploits the fact that the subsets of size  $k$  of the set  $\{1, \dots, k+2\}$  are uniquely identified by their pairwise intersection sizes.

### 4. Determining the private images

Suppose the attacker has identified a floral matrix in the previous steps, which corresponds to the selection vectors  $w_{i1}, \dots, w_{it}$ , where  $t = \binom{k+2}{k}$ . The structure of the floral matrix encodes information about the indices of private images that are common between pairs of selection vectors. Specifically, the row and column indices of the matrix indicate which private images are shared. This allows the attacker to construct a system of equations of the form  $|\langle w_{ij}, p_l \rangle| = \tilde{x}_l$  for all  $l \in 1, \dots, d$ , where  $p_l$  denotes the private images and  $\tilde{x}_l$  are known quantities.

From another perspective, each row or column of the floral submatrix can be indexed by a subset of size  $k$  from a set of size  $k+2$ . Each element in such a subset represents the index of a private image. For any given element in the floral matrix—which itself is a submatrix of the Gram matrix  $M$ —the position of the element along the rows provides the attacker with a set of  $k$  private image indices, while the position along the columns provides another set of  $k$  indices. By intersecting these two sets, the attacker can determine which private images are common between the selection vectors associated with the corresponding row and column. Solving the resulting system of equations enables the attacker to recover the indices of the private images.

## B LUO ATTACK

In Luo et al. (2022), the authors observed that the method proposed by Carlini in Carlini et al. (2021a) can be mitigated by applying data augmentation before the mixup process. To address this, they introduce a new approach that successfully bypasses this mitigation strategy. Their method operates as follows:

1. In the first step, the attacker computes the absolute value of each pixel in every encoded image.

2. Next, a similarity score is calculated for every pair of encoded images to determine, with high probability, whether a given pair is derived from the same private image. To compute this score, the authors propose a comparative network that takes as input both high-resolution and low-resolution versions of the image pairs. This approach yields better results than the standard ResNet architecture used by Carlini. Based on the similarity scores,

702 the attacker clusters the encoded images, with each cluster corresponding to a distinct pri-  
 703 vate image.  
 704

- 705 For each cluster obtained in the previous step, the attacker re-weights all encoded images  
 706 using the weights associated with the corresponding private image. These weights can  
 707 be easily inferred from the associated encoded labels. Subsequently, a neural network  
 708 is trained to perform image relaxation and fusion. This strategy counteracts the effects of  
 709 geometric image augmentation by generating a set of features that are invariant to geometric  
 710 transformations. An initial version of the private image is then constructed in the fusion  
 711 step by combining these feature maps.
- 712 In the final step, the attacker trains an additional neural network to denoise the image pro-  
 713 duced in the previous stage.

714 **C CARLINI ATTACK**

715 The attack consists of two main stages. In the first stage, the attacker determines the two private  
 716 images used to generate each encoded image during the mixing process. In the second stage, the  
 717 attacker reconstructs the private images by solving a noisy linear system of equations:  
 718

- 720 The attacker computes the absolute value of each mixup encoding to counteract the random  
 721 sign changes introduced by the mask  $\sigma_i$  in (3).
- 722 To identify whether two encoded images share at least one common private image, the at-  
 723 tacker calculates a similarity function between each pair of encoded images. This similarity  
 724 function is approximated using a neural network trained on public data transformed via the  
 725 mixup algorithm. Using the similarity scores, the attacker constructs a weighted graph  
 726 where vertices represent encoded samples, and edge weights correspond to the similarity  
 727 function's output.
- 728 Based on the weighted graph, the attacker identifies densely connected cliques, enabling  
 729 clustering of encoded samples that share a common private image. Each cluster is repre-  
 730 sented as a set  $S_i$ ,  $1 \leq i \leq n$ , where each set contains encoded samples derived from the  
 731 same private image.
- 732 Since each encoded image is generated by mixing two private images, the attacker con-  
 733 structs a bipartite similarity graph connecting encoded images to the sets identified in the  
 734 previous step. Edge weights represent the distance between an encoded image  $x_i$  and a set  
 735  $S_i$ . This step determines, for each encoded image, the two sets corresponding to the private  
 736 images used in its construction.
- 737 Using the bipartite graph, the attacker maps each encoded image to two sets, representing  
 738 the private images involved in its generation during the mixup process.
- 739 The attacker recovers the weights used to generate each encoded image by analyzing the  
 740 mixup of the labels, as described in (4). Since the labels are one-hot encoded, recovering  
 741 the associated weights is straightforward.
- 742 Finally, the attacker constructs a matrix  $B \in \mathbb{R}^{n \times d}$ , where each row corresponds to an  
 743 encoded image  $\tilde{x}_i$ , i.e.,  $B_i = \tilde{x}_i$ . A sparse matrix  $M \in \mathbb{R}^{n \times n}$  is also constructed, where  
 744 each row contains two non-zero entries representing the weights  $w_1^i$  and  $w_2^i$  associated  
 745 with the private images used to compute the corresponding encoded image. Let  $A \in \mathbb{R}^{n \times d}$   
 746 represent the matrix of private images, where each row  $A_i = x_i$ . The attacker solves the  
 747 noisy linear system  $B = M \cdot A + e$ , where  $e$  represents the public images used in the mixup.  
 748 This system can be efficiently solved using gradient descent.

749 **D MORE RELATED WORK**

750 Liu et al. (2020) proposed a different approach, where a classifier is trained on mixup samples and  
 751 images to produce mixup results that can later be de-mixed. Unlike previous methods, this approach  
 752 does not involve training on mixup samples followed by inference on original data. Instead, both  
 753 training and inference are performed on mixup data, with the inference process generating mixup  
 754 results that can then be used to recover the correct labels.  
 755

756 In a more recent study, Wang et al. (2024) proposed a mixup-like approach to mitigate model inver-  
 757 sion attacks on face recognition systems. Instead of mixing images directly, the authors suggested  
 758 mixing samples in the frequency domain. Additionally, they employed a reinforcement learning  
 759 strategy to dynamically determine the number of images to mix, balancing privacy and utility. Simi-  
 760 larly, Xiang et al. (2023) introduced a mixing strategy to preserve image privacy during training.  
 761 Their method involves splitting each image into multiple blocks and replacing parts of these blocks  
 762 with corresponding blocks from other images with the same label. In another study, Li et al. pro-  
 763 posed a new privacy metric called Visual Feature Entropy (VFE), calculated for a region of an image  
 764 as the sum of squared gradients with respect to both axes. This metric aims to quantify the amount  
 765 of information that needs protection by analyzing the entropy of a region. The authors' mixing strat-  
 766 egy involves shuffling pixels within an image based on the VFE metric. Although this method does  
 767 not involve computing a weighted sum, it can be interpreted as a form of intra-image data mixing.  
 768 Eloul et al. (2024) present the concept of mixing gradients in federated learning to enhance security  
 769 against gradient inversion attacks. Although their method does not involve using random weights for  
 770 gradient mixing, their straightforward approach of directly averaging gradients across a batch, com-  
 771 bined with modifications to the loss function, significantly improves resistance to gradient inversion  
 772 attacks.

772 The concept of data mixing is rooted in the broader idea of learnable obfuscation, which encom-  
 773 passes techniques designed to transform data in a way that allows algorithms to learn from the  
 774 transformed data while safeguarding the privacy of the original data He et al. (2020); Yala et al.;  
 775 Taki & Mastorakis (2024); Popescu et al. (2022); Nythia et al. (2017). For instance, in Nythia et al.  
 776 (2017), the authors propose using the Arnold transformation to scramble images before inputting  
 777 them into a face recognition system. This transformation rearranges image pixels by mapping each  
 778 pixel to a new location determined by a linear transformation.

779 In Popescu et al. (2022), a method combining Variational Autoencoders (VAEs) with a substitution  
 780 technique is introduced to protect medical images during neural network analysis. The approach  
 781 involves training a VAE to reconstruct the image and then applying a substitution table to the latent  
 782 space representation of the data. Similarly, Taki & Mastorakis (2024) presents a method to ensure the  
 783 privacy of both training data and neural network architecture. For image data, the authors propose  
 784 transforming it into a higher-dimensional space. To protect the architecture, they introduce random  
 785 subnetworks with synthetic parameters that do not affect the network's accuracy or data flow.

786 The NeuraCrypt method, proposed in Yala et al., protects data privacy by transforming it with a ran-  
 787 dom neural network. This approach is extended to enable privacy-preserving collaborative training,  
 788 where all parties share transformed data with a central server. For the server to learn patterns from  
 789 the combined datasets, all parties must use the same neural network for data transformation. Fi-  
 790 nally, He et al. (2020) introduces a privacy-preserving method that applies a linear transformation to  
 791 each data sample. The authors also provide formal proofs demonstrating the information-theoretic  
 792 security of their approach under specific conditions.

793 A common characteristic of learnable obfuscation techniques is that the same transforma-  
 794 tion—though potentially generated using independently chosen random parameters—must be ap-  
 795 plied to all samples in the dataset being protected. This creates a notable vulnerability: such tech-  
 796 niques cannot provide security against chosen-plaintext attacks. This limitation, formally introduced  
 797 and proven in Xiao et al. (2024), highlights an inherent weakness in these methods. Informally,  
 798 learnable obfuscation can protect the privacy of plaintext data only under the assumption that the  
 799 attacker does not have prior knowledge of the original data.

800 At first glance, this assumption may seem reasonable, as protecting data already known to an at-  
 801 tacker might appear unnecessary. However, in practical scenarios, this assumption often fails. For  
 802 instance, to improve the generalization capabilities of a machine learning model, private datasets  
 803 are frequently augmented with publicly available data. For example, a private image dataset might  
 804 be enriched with images from CIFAR-100 Krizhevsky et al. (2009). To preserve the discriminative  
 805 properties of the data and enable the model to generalize, the added public data must undergo the  
 806 same transformation used to protect the private dataset.

807 This practice introduces a significant risk: an attacker with access to both the original public dataset  
 808 and its transformed version could potentially design algorithms to reverse-engineer the transfor-  
 809 mation applied to the private data. An example of this vulnerability is described in Carlini et al.

(2021b), where the authors successfully developed an algorithm to solve the NeuraCrypt challenge Yala et al., effectively bypassing the intended privacy protections.

## E EXPERIMENTAL DETAILS

All experiments were developed using the PyTorch framework and performed on an NVIDIA L4 GPU with 24 GB of available VRAM. Across all datasets, we consistently used a batch size of 128. For optimization, we employed the AdamW optimizer with a weight decay of  $1 \times 10^{-4}$ . The initial learning rate was set to 0.001 for all benchmarks, and we utilized a cosine annealing learning rate scheduler.

Our experimental evaluation was conducted on three distinct benchmarks using ResNet architectures (He et al., 2016), with specific configurations detailed in Table 5. A key aspect of our methodology is the exclusive use of the feature extraction layers from these architectures; the final classifier layers were omitted as our focus is on training features. For all datasets we apply similar transformations, which consist of a resize operation (224 pixels height and width) and a normalization. The MNIST dataset is also adjusted such that it has three channels, making it compatible with the chosen architectures.

Table 5: The configurations used within experiments for each dataset. The mean and standard deviation values for MNIST are for a single channel, while for CIFAR datasets they correspond to the (R, G, B) channels.

Dataset	Architecture	Epochs	Mean	Std.
MNIST	ResNet18	120	0.1307	0.3081
CIFAR-10	ResNet18	200	[0.4914, 0.4822, 0.4465]	[0.2470, 0.2435, 0.2616]
CIFAR-100	ResNet50	200	[0.5071, 0.4867, 0.4408]	[0.2675, 0.2565, 0.2761]
Tiny-ImageNet	ResNet50	200	[0.5071, 0.4867, 0.4408]	[0.2675, 0.2565, 0.2761]

The training objective was to minimize the following loss function, which is designed for our Mixup implementation:

$$\mathcal{L} = -\frac{1}{N} \sum_{i=1}^N \sum_{j=1}^C y_{ij} \log(\text{softmax}(\mathbf{o}_i)_j) \quad (7)$$

where  $N$  is the batch size,  $C$  is the number of classes,  $\mathbf{y}_i$  is the (potentially soft) label for sample  $i$ , and  $\mathbf{o}_i$  is the model output for sample  $i$ . In PyTorch, this is implemented as:

```
loss = -(labels * torch.log_softmax(outputs, dim=1)).sum(dim=1).mean()
```

For the final classification step, we used a custom feed-forward neural network with three dense layers. The first and second of these dense layers are followed by batch normalization, GELU activation, and then a dropout. The precise structure of this neural network is described by the following equation:

$$\begin{aligned} \text{Classifier}(x; n_{\text{cls}}) &= \text{Flatten}(x) \rightarrow \text{Linear}_{in; 1024} \rightarrow \text{BN}_{1024} \rightarrow \text{GELU} \rightarrow \text{Dropout}(0.5) \\ &\rightarrow \text{Linear}_{1024; 512} \rightarrow \text{BN}_{512} \rightarrow \text{GELU} \rightarrow \text{Dropout}(0.5) \\ &\rightarrow \text{Linear}_{512; n_{\text{cls}}} \end{aligned} \quad (8)$$

The  $n_{\text{cls}}$  term represents the number of classes that the classification must be made on (e.g., 10 for MNIST). The  $in$  dimension of the flattened tensor  $x$  within the first linear layer is 25088 ( $512 \times 7 \times 7$  for ResNet18 and ResNet34).

## F ATTACK IMPLEMENTATION DETAILS

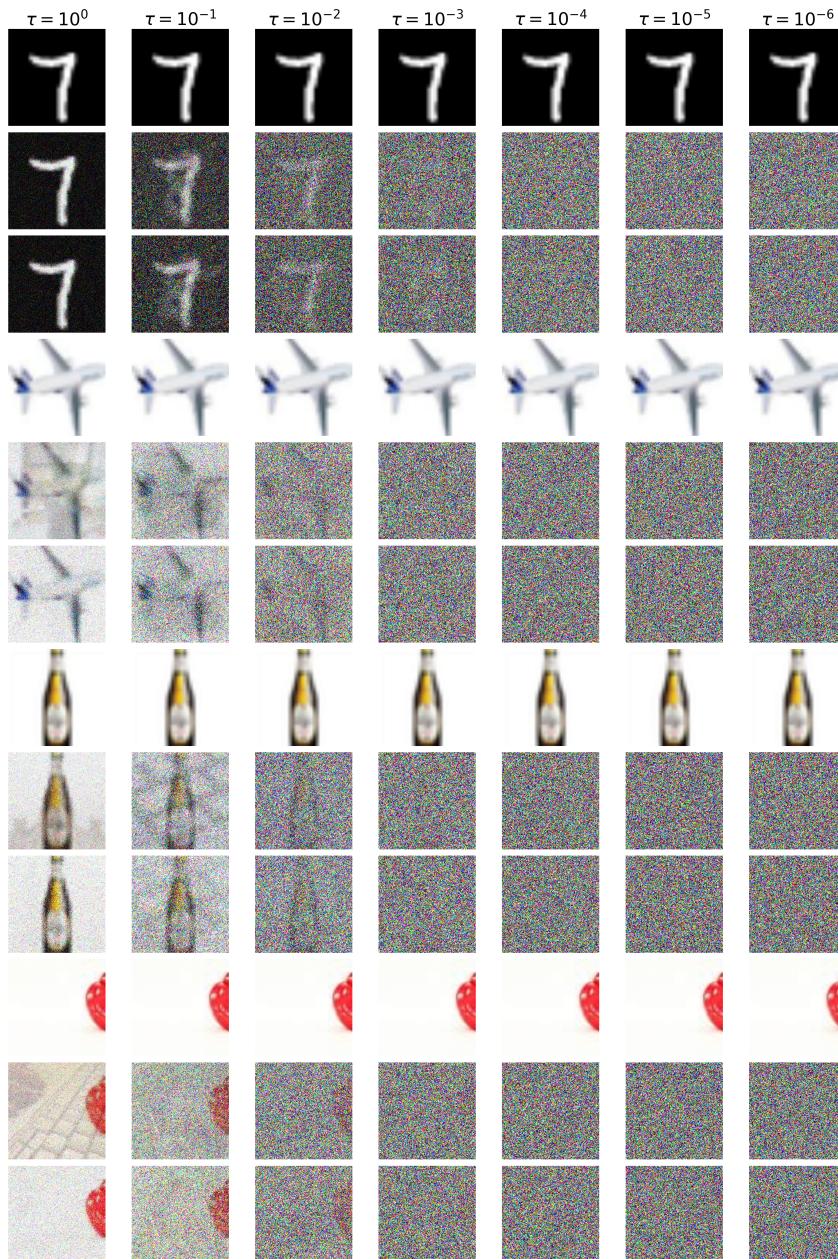
All attacks use the same preprocessing described in Sec. E, including resizing inputs to  $224 \times 224$ , ImageNet-style normalization, and converting MNIST to three channels. Experiments are executed on a single NVIDIA L4 GPU (24 GB) with fixed random seed (0).

864 For the *linear reconstruction attack*, we optimize a recovery tensor using the PyTorch Adam opti-  
 865 mizer for 200 steps with learning rate 0.05. We include anisotropic total variation regularization and  
 866 an  $\ell_2$  penalty with default weights  $\lambda_{\text{tv}} = 10^{-3}$  and  $\lambda_2 = 10^{-4}$ . After each step, recovered tensors  
 867 are clamped back into the normalized image range. Recovery quality is measured using per-image  
 868 SNR in dB, reported as mean  $\pm$  std. A single visualization panel is produced: columns correspond  
 869 to different  $\tau$  values and, for each dataset, three aligned rows display the original, the Mixup input,  
 870 and the recovered image. The displayed index per dataset is chosen as the best-recovered example  
 871 at the largest  $\tau$ .

872 For the *non-linear attack*, we train a U-Net denoiser on Tiny-ImageNet Mixup pairs and eval-  
 873 uate zero-shot on CIFAR-10. The U-Net follows a three-level encoder-decoder design with base  
 874 width  $B = 48$ . The encoder consists of successive DoubleConv blocks with channel pro-  
 875 gression ( $3 \rightarrow B \rightarrow 2B \rightarrow 4B \rightarrow 8B$ ), separated by  $2 \times 2$  max-pooling. The decoder mirrors this structure  
 876 with transposed convolutions for upsampling, skip connections from encoder features, DoubleConv  
 877 blocks with channel progression ( $8B \rightarrow 4B \rightarrow 2B \rightarrow B$ ), and a final  $1 \times 1$  convolution mapping to three  
 878 output channels. All layers use ReLU activations, and outputs are clamped to the normalized range.  
 879 Training uses Adam with batch size 32, learning rate  $10^{-3}$ , 30 epochs, and an  $\ell_1$  loss augmented  
 880 with a small TV penalty ( $10^{-4}$ ). Partner selection is deterministic per index so the same samples  
 881 align across  $\tau$  values. Evaluation again reports SNR (mean  $\pm$  std). A single summary panel is  
 882 generated: the first row repeats the ground-truth image associated with the best reconstruction at the  
 883 largest  $\tau$ , while the second and third rows show the corresponding Mixup and recovered images for  
 884 each  $\tau$  (displayed as columns). Only the  $\tau$  labels appear above columns.

885 All runs save a single figure per attack configuration along with a timestamped log containing the  
 886 full console output.

887  
 888  
 889  
 890  
 891  
 892  
 893  
 894  
 895  
 896  
 897  
 898  
 899  
 900  
 901  
 902  
 903  
 904  
 905  
 906  
 907  
 908  
 909  
 910  
 911  
 912  
 913  
 914  
 915  
 916  
 917

918 G FULL FIGURE 3  
919  
920  
921962 Figure 3: Recovered images under the linear attack for decreasing noise levels  $\tau$   
963964 H THE USE OF LARGE LANGUAGE MODELS  
965

966 We utilized Large Language Models (LLMs) in three specific ways during this work. First, after  
967 conducting a manual review of the state of the art using traditional search engines such as Google  
968 Scholar, we used LLMs to assist in identifying additional relevant papers. Second, LLMs were  
969 employed to help implement the experiments described in this study. Third, LLMs were used for  
970 grammar correction and minor improvements to the flow of the text. Importantly, LLMs were not  
971 used to generate or write any paragraphs; their role in writing was limited to minor edits and en-  
hancements.