
Replication study of "Privacy-preserving Collaborative Learning with Automatic Transformation Search"

Anonymous Author(s)

Affiliation

Address

email

Reproducibility Summary

1

2 **Scope of Reproducibility**

3 We evaluate the reproducibility of this paper, which proposes an automatic search algorithm to find privacy preserving
4 transformation policies in the setting of federated learning. To achieve this we test all the main claims made by the
5 authors by rerunning the experiments and reporting the reproduced results. We further extend their work to a new
6 dataset.

7 **Methodology**

8 We perform all experiments using the model architectures and hyperparameters proposed by the authors. We use
9 the same datasets and extend their work to include one new dataset. A codebase was available which enables us to
10 reproduce some of the results. However we deliver a contribution by fully re-implementing the codebase in PyTorch
11 Lightning to ensure all components are modular, and experiments can be easily executed and extended, to the benefit of
12 future research using the authors' method. All experiments are performed on Nvidia GTX 1080 GPUs.

13 **Results**

14 Overall we find the same results as the authors: searched transformation policies can defend users in federated learning
15 from reconstruction attacks. These transformations also have negligible impact on training efficiency and model
16 accuracy. However we do not observe the reported correlation between the authors privacy-score and PSNR. We are
17 in contact with the authors about this. Also we find that the results differ greatly from image to image, with standard
18 deviations in PSNR values of over 25% the value. This means that for some specific images the method is not effective.

19 **What was easy**

20 Paper was clearly written and the general idea was easy to follow. There was a codebase available in PyTorch and part
21 of the experiments were reproducible using this code.

22 **What was difficult**

23 The codebase was not clearly structured and has to be altered to produce results for most experiments reported in the
24 paper. The reimplementing of the codebase was non-trivial due to otherwise undocumented details in the code having
25 a large impact on outcomes.

26 **Communication with original authors**

27 The authors were contacted on multiple issues regarding implementation details and notation in the paper. Most of
28 these were resolved swiftly and constructively. On two issues we remain in contact with the authors at this time.

29 1 Introduction

30 Collaborative learning systems allow multiple users to jointly train a Deep Learning (DL) model. Each user has their
31 own training data which is to calculate local gradients [18][7][12]. These local gradients are then shared among all
32 users to update the parameters of the shared DL model, without the need of sensitive data leaving user’s device. Primary
33 benefit of federated learning is its capacity to improve the generalization of the resulting model, while maintaining
34 privacy over the training data of individual users. This is especially important as confidentiality quickly becomes an
35 essential quality of DL models [1]. Because of that, federated learning is used in applications from mobile networks
36 [10] to autonomous driving [13] and health care [2].

37 However, this privacy benefit can be undone by reconstruction attacks as proposed by [6][19][20]. These attacks make it
38 possible to reconstruct the original private training samples of users from the shared gradients of the federated learning
39 system. This poses a considerable threat to the privacy of users of federated learning systems and the confidentiality of
40 their data samples.

41 The paper subject to this reproducibility study proposes a novel approach to mitigate the threat from reconstruction
42 attacks by augmenting the local training data of the user, before calculating the gradients [5]. Furthermore, the authors
43 develop an automatic search algorithm to find the optimal transformation policies to augment the data and propose two
44 novel metrics, S_{pri} and S_{acc} , to increase the efficiency of this search.

45 In this reproducibility report, we evaluate the main claims made by the authors of [5] by reproducing their experiments.
46 Moreover, we assess the availability of hyperparameters and other information needed for reproducibility, as well as
47 discuss the usability of the provided codebase. We also extend the experimental setup towards a new dataset.

48 2 Scope of reproducibility

49 The main goal of the original paper is to develop an automatic search algorithm to find transformation policies that can
50 defend privacy-sensitive training data against reconstruction attacks in a federated learning system. To achieve this,
51 authors devise two novel metrics, described in Section 3.2. The main claims made in the paper are the following:

- 52 • **Claim 1:** by augmenting training samples with carefully-selected transformation policies, reconstruction
53 attacks become infeasible
- 54 • **Claim 2:** the proposed search algorithm can find good and general policies, i.e. policies that are able to defeat
55 multiple variants of reconstruction attacks
- 56 • **Claim 3:** the found policies are highly transferable; good policies searched for one dataset are also suitable for
57 another datasets
- 58 • **Claim 4:** the found policies have negligible impact on the training efficiency
- 59 • **Claim 5:** in general, a good policy is made up of transformations that distort the details of the training samples,
60 while maintaining the semantic information
- 61 • **Claim 6:** the five transformations that work best are *horizontal shifting (9)*, *brightness (9)*, *brightness (6)*,
62 *contrast (7)* and *contrast (6)* (number inside the brackets represents the intensity of the applied transformation)
- 63 • **Claim 7:** S_{pri} is a good measure of privacy; it is linearly correlated to Peak signal-to-noise ratio (PSNR) [9]
64 with a Pearson Coefficient [15] of 0.697

65 Each of these claims is supported by the results of one or more experiments in [5], represented in the tables and figures.
66 In this reproducibility study, we rerun the experiments and reproduce the resulting tables and figures. In Section 5,
67 we list which experiments support which claims. In Section 6, we discuss the reproducibility of each experiment and
68 evaluate the validity of the claims.

69 Beyond reproducing the above claims from the original paper, we propose two extensions. Both of these extensions are
70 based on the transferability of the searched policies as claimed in *Claim 3*. We test the transferability of the policies
71 against additional dataset and evaluate whether the best performing transformations are the same on this dataset.

72 **Extension 1:** Using the policies searched on one dataset and applying them to a new dataset can make reconstruction
73 attacks against this new dataset infeasible

74 **Extension 2:** Since good policies share the same general qualities, as claimed by *Claim 5*, the five best transformations
75 from *Claim 6* are the same when using a different dataset.

76 In Section 5, we show the results for these extensions, and in Section 6, we relate them to the claims, experiments, and
77 results from the original paper.

78 **3 Finding privacy-preserving transformation policies**

79 The original paper proposes an automatic search algorithm for finding privacy-preserving transformation policies. To
80 better understand this main contribution, we take an in-depth look at what a transformation policy is and how good
81 policies are found within a reasonable time.

82 **3.1 Transformation policies**

83 Transformations or augmentations have been widely used to improve model performance and generalizability in DL.
84 In [5], transformations from AutoAugment¹ [3] are repurposed to protect sensitive training data from reconstruction
85 attacks. The library contains 50 different transformations, including rotation, crop, shift, inversion, brightness, and
86 contrast. A *transformation policy* is a combination of k such transformations applied to the training samples. In [5],
87 $k = 3$ is chosen and the policies are denoted by the indices of the transformations within the AutoAugment library.

88 Consistently apply the best policy to the data would risk domain shift in the dataset. Therefore, the authors propose the
89 *hybrid strategy*, where a policy is randomly selected from the candidate policies - this way, good privacy and accuracy
90 are guaranteed [5].

91 **3.2 Reducing the search-space**

92 To find candidate policies, it is necessary to determine their effect on both privacy and accuracy. The transformations
93 must be applied to training data, and a model must be trained. Because fully training a model is very expensive, the
94 authors propose two metrics that serve as a proxy for the privacy preservation and accuracy of the fully trained model:
95 $\text{privacy-score}(S_{pri})$ and $\text{accuracy-score}(S_{acc})$. Low S_{pri} entails the model has high privacy preservation potential,
96 whereas high S_{acc} means the model achieves good accuracy with the applied transformation policies. These metrics
97 produce results on model that are trained with only 10% of the data for only 25% training iterations, reducing the
98 search-space and making the policy search feasible in a reasonable time. Further details about the definition of S_{pri} and
99 S_{acc} can be found in sections 4.2 and 4.3 of [5].

100 **4 Experimental setup and code**

101 To verify the claims made by the authors of [5], we reproduce their experiments. These experiments roughly fall
102 into four categories: evaluating the effectiveness of the searched policies against reconstruction attacks, testing the
103 transferability of the searched policies on different datasets and models, checking the impact on model efficiency,
104 and studying the semantics behind the different transformations. Multiple models must be trained on augmented and
105 un-augmented data for all these categories. For the attacks, the approach from [6] is applied. Section 5 provides a
106 detailed description of the experiments and shows the results.

107 To reproduce the experiments performed by the authors, we used their existing codebase², which is implemented in
108 PyTorch [14]. We refactored parts of this code and re-implemented the rest to our own version written in PyTorch
109 Lightning³, which leverages the interface advantages of the Lightning framework to make running experiments and
110 logging results more intuitive. Main benefit of doing so is that more experiments can be tested with finer clarity
111 and control of the setup. Our refactoring is this study's main contribution, and the codebase is publicly available at
112 <https://anonymous.4open.science/r/MLRC2021-0454>.

¹<https://github.com/DeepVoltaire/AutoAugment>

²<https://github.com/gaow0007/ATSPrivacy>

³<https://github.com/PyTorchLightning/pytorch-lightning>

113 4.1 Datasets

114 The experiments in [5] are performed on two datasets, CIFAR-100⁴ [11], and Fashion-MNIST⁵ [17]. CIFAR-100
115 contains 60,000 color images of size 32×32 , from 100 classes. The test set is used as the validation set, consistent
116 with the authors' codebase. On the other hand, the Fashion-MNIST dataset contains 70,000 grey-scale images of
117 28×28 resolution from 10 classes. Again the test set is used as a validation set. We run experiments on one additional
118 dataset in our extensions - Tiny ImageNet200⁶ [4]. It contains 120,000, 64×64 RGB images of 200 different classes.
119 However, a *tiny* version of the dataset is introduced in the original paper for policy-search purposes. This dataset version
120 contains 10% of the original samples, using the same distribution. It's later used to train the models for the evaluation
121 of S_{pri} and S_{acc} in the search algorithm.

122 4.2 Model descriptions

123 We use the following models:

- 124 • ResNet20-4, a variation of ResNet20 [8] that has four times the number of channels also used in [6]. The total
125 number of parameters is 4.4M.
- 126 • ConvNet [6] - 8-layer Convolutional Neural Network, with batch normalization and a ReLU layer after each
127 convolution layer. For this model the total number of parameters is 3.7M.

128 The original codebase uses the implementation of both models from the repository⁷ of [6]. Our models are re-
129 implemented in Pytorch Lightning. Both models were compared with the models from the original codebase in terms of
130 accuracy; they achieved comparable results.

131 4.3 Hyperparameters

132 For policy search, we used $C_{max} = 1500$ and *max policies* equal to 10. The batch size was 128, and the number of
133 transforms in policy was 3. For training, the batch size was also 128 and the number of epochs was 60 (see Section
134 4.4). To obtain a semi-trained network, we used a subset of 10% of the training dataset. The attack is performed on
135 image with index 0, and we reused the remaining setups according to the original paper e.g. "inversed" (default attack).
136 Except for Figure 4, where a default config was used, with a number of maximum iterations changed to 2500. For
137 further experiments, we followed the same conventions.

138 4.4 Computational requirements

139 We ran our experiments using Nvidia GeForce GTX 1080 GPU. The policy search took approximately 10 hours. The
140 training of one model took approximately 2h 40min using the original approach. However, training for 60 epochs
141 achieves the same accuracy but in 50 minutes. It is because there exist periods of plateau, while lr is not scheduled yet
142 to drop. One attack with 2500 iterations took approximately 5 minutes, so measuring the correlation between S_{pri} and
143 PSNR took 8.5 hours (with policy search).

144 5 Experiments and results

145 5.1 Results reproducing original paper

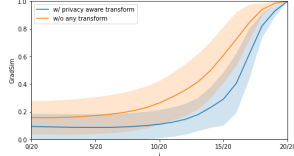
146 **Experiment 1** A reconstruction attack on 100 images from the CIFAR-100 validation set is performed with and
147 without a searched transformation policy applied. We document the optimization process of the attack in terms of
148 GradSim. The model used is ResNet20 trained on the tiny dataset for 50 epochs. The results of this experiment are
149 shown in Figure 2, which shows a very similar result to the original paper. In addition to the original figure, we show the
150 standard deviation over the 100 images, since GradSim can differ significantly from image to image. When taking the
151 average of multiple runs, it can be seen that the privacy-aware transform does indeed make the GradSim convergence
152 more difficult.

⁴<https://www.cs.toronto.edu/~kriz/cifar.html>

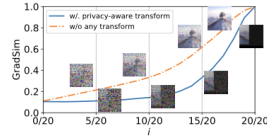
⁵<https://github.com/zalandoresearch/fashion-mnist>

⁶<http://cs231n.stanford.edu/tiny-imagenet-200.zip>

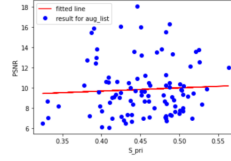
⁷<https://github.com/JonasGeiping/invertinggradients>



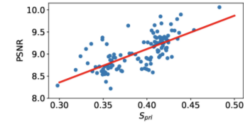
(a) Reproduced result



(b) Result from original paper



(a) Reproduced results



(b) Original results

Figure 2: Optimization process of reconstruction attack with and without searched policy

Figure 4: Correlation between S_{pri} and PSNR

153 **Experiment 2** A visual comparison between reconstructed images with and without a searched transformation policy
 154 applied is performed for both ResNet20 and ConvNet on images from CIFAR-100 and Fashion-MNIST. The optimizer
 155 used in the attack is Adam+Cosine. The images, the resulting reconstructions, and their PSNR values are shown in the
 156 left half of Figure 5. The results from the original paper are shown at the right side of Figure 5. As can be seen, the
 157 images used and PSNR values reported are different. This is due to the fact that it was too expensive to identify the
 158 exact same images and PSNR values differ quite severely depending on the image used. However, for all 12 images, we
 159 observe a less pronounced visual effect of the transformation policy as well as a smaller gap in PSNR values between
 160 the reconstructions with and without the policies applied. This implicates that the effect shown in the original paper is
 not as severe for all images, although the images we selected may be particularly easy to reconstruct.




Figure 5: Visualization results for reconstruction attacks on different datasets and models with associated PSNR values. Our results above and original results below.

161

162 **Experiment 3** To gain further insight into the effectiveness of the different policies, we report the qualitative and
 163 quantitative results of Adam+Cosine attacks and model accuracy for the datasets and models in Figure 5. The results
 164 are calculated over 6 images as performing the experiment is very expensive and number wasn't stated in the paper. The
 policies considered and the results are listed in Table 1. Table 1 shows similar patterns to the original paper, where the

Policy	PSNR	PSNR (std)	Acc
None	12.15	2.06	78.11
Random	9.92	1.93	75.02
3-1-7	6.77	0.88	71.59
43-18-18	9.34	1.81	77.16
Hybrid	8.25	1.64	77.47

(a) CIFAR-100 + ResNet20

Policy	PSNR	PSNR (std)	Acc
None	9.81	4.41	95.19
Random	10.06	2.04	95.19
19-15-45	8.26	0.37	92.44
2-43-21	8.93	2.93	93.93
Hybrid	8.41	1.45	95.14

(c) FMNIST + ResNet20

Policy	PSNR	PSNR (std)	Acc
None	11.44	2.93	72.97
Random	10.29	1.02	71.93
21-13-3	8.23	2.18	63.26
7-4-15	10.31	2.14	70.77
Hybrid	9.89	1.47	68.91

(b) CIFAR-100 + ConvNet

Policy	PSNR	PSNR (std)	Acc
None	9.52	3.27	94.61
Random	9.47	2.27	94.47
42-28-42	7.59	0.89	94.62
14-48-48	8.41	2.10	94.68
Hybrid	6.80	0.98	94.59

(d) FMNIST + ConvNet

Table 1: PSNR (db) (including mean and standard deviation over 6 images) and model accuracy (%) of different transformation configurations for each model and dataset. 19 – 1 – 18 is the random policy.

165 searched policies have low PSNR values compared to not using transformations. We do observe that PSNR values have
 166 a relatively high standard deviation, and during our experiments, we found that the policies do not form a good defense
 167 for some images. This problem will be further discussed in Section 6.

168 **Experiment 4** The defensive qualities of the searched transformation policies are benchmarked against existing
 169 defenses from the literature [20] [16] under the Adam+Cosine attack. The results are shown in Table 6. Although the
 170 exact values differ slightly, the overall results are similar to the original paper, where all the existing defenses perform
 171 worse than the hybrid strategy.

172 **Experiment 5** This experiment concerns **Claim 2**. Because policies should be general, they are tested against various
 173 attack configurations. For this, we again use 6 images from the test set and perform the different attacks on the images
 174 without the transformation policies applied and with the hybrid strategy transformation policies applied. The results
 175 are shown in Table 2. As can be seen from the table, the hybrid strategy works well against all configurations of the
 reconstruction attack. This is in line with the results from the original paper.

Attack	None	None (std)	Hybrid	Hybrid (std)
LBFGS+L2	8.61	1.22	6.33	2.00
Adam+Cosine	12.15	2.06	8.25	1.64
LBFGS+Cosine	9.62	0.91	7.47	0.25
Adam+L1	9.48	0.71	6.43	0.16
Adam+L2	9.28	0.69	6.46	0.21
SGD+Cosine	12.60	2.07	8.03	1.47

Table 2: PSNR values (db) (including mean and standard deviation over 6 images) of reconstructed images with and without transformations applied for different attack configurations

Policy	PSNR	PSNR std
None	15.39	2.78
3-1-7	8.47	0.85
43-18-18	10.97	1.06
Hybrid	8.95	0.90

Table 3: CIFAR100 with ResNet20

177 **Experiment 6** This experiment concerns the transferability of **Claim 3**. To test this, the policies searched on CIFAR-
 178 100 are applied to Fashion-MNIST using both ResNet20 and ConvNet. Reconstruction attacks are performed with
 179 the Adam+Cosine attack. The resulting PSNR values and accuracies are listed in Table 4. The results differ from the
 original. It can be seen that the transformation policies are not effective here.

Policy	PSNR	PSNR (std)	Acc
None	9.81	4.41	95.19
3-1-7	9.30	2.72	93.20
43-18-18	10.03	2.23	94.88
Hybrid	7.49	1.57	94.49

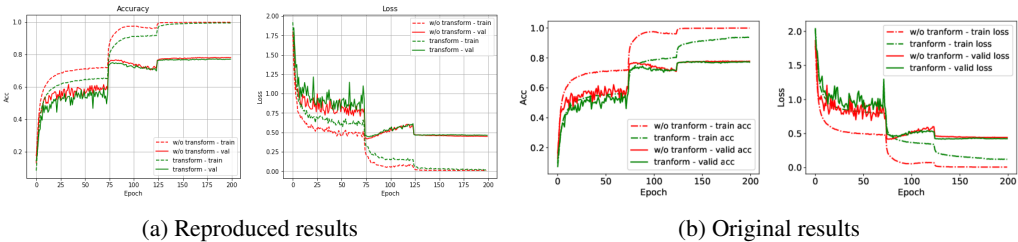
(a) FMNIST + ResNet20

Policy	PSNR	PSNR (std)	Acc
None	9.52	3.27	94.61
21-13-3	9.99	2.12	92.38
7-4-15	9.34	1.62	94.35
Hybrid	11.50	5.80	93.77

(b) FMNIST + ConvNet

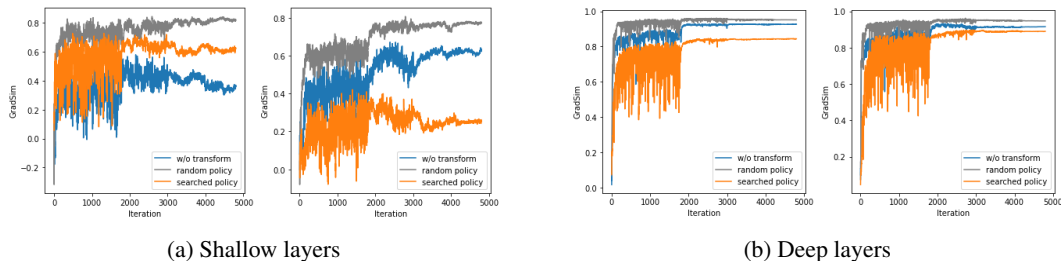
180 Table 4: Resulting PSNR (dB) and accuracy (%) values for applying policies searched on CIFAR-100 to Fashion-MINST

181 **Experiment 7** The following experiment is aimed at **Claim 4**. The authors state that applying the search policies has
 182 a negligible impact on training efficiency. We trained ResNet20 with the searched policies applied and documented the
 183 loss and accuracy convergence to test this. From Figure 6 it can be seen that indeed applying transformations has almost
 184 zero impact on the training efficiency. It is also noteworthy to observe that the training curves are almost identical
 compared with the results from the original work.



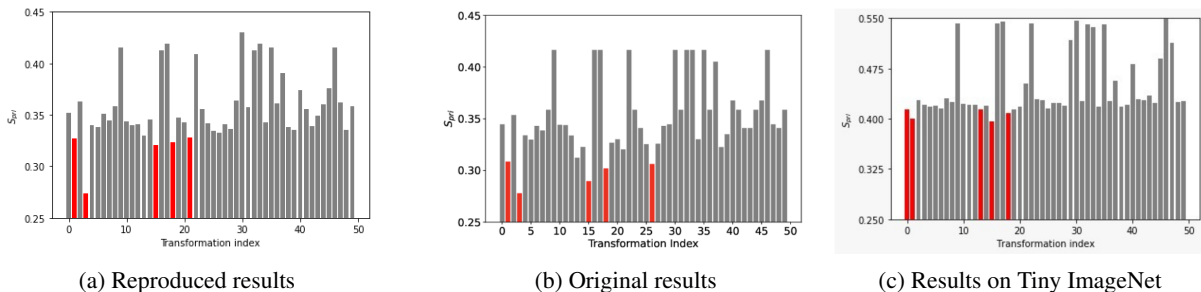
185 Figure 6: Convergence speed with and without transformations applied

186 **Experiment 8** **Claim 5** states that good transformation policies obfuscate details in the training samples but maintain
 187 high-order semantic information. As such, attackers will have trouble reconstructing high frequency information. We
 188 test this by comparing the attacker-defender gradient similarity during an attack of models trained with the searched
 189 policy, a random policy, and no policy applied. From Figure 7, it can be seen that in shallow layers, the gradients differ
 190 significantly, whereas in deep layers, the gradients are very similar. This implies that the transformations do indeed
 have the desired effect and is in line with the results from the original paper.



191 Figure 7: Reproduced results of gradient similarity during the reconstruction optimization, for CIFAR100 with ResNet20

192 **Experiment 9** In **Claim 6** the authors report their 5 top transformations. We test whether we can find the same ones
 193 by calculating the privacy score on the dataset for each individual augmentation and show the results in Figure 8a and
 8b. Out of the best 5 transformations reported in the original paper we found 4 overlapping ones.



194 Figure 8: Privacy scores of the 50 transformation functions in the augmentation library, best transformations are red.

195 **Experiment 10** The final experiment reproducing the results from the original paper is aimed at **Claim 7**. The authors
 196 claim that their privacy-score S_{pri} is linearly correlated with PSNR with a Pearson-coefficient of 0.697. We test this
 197 by running attacks and evaluating S_{pri} on the model trained on tiny cifar100 for 50 epochs and found a very different
 198 result. As shown in Figure 4 there is hardly any correlation (Pearson-coefficient is 0.123). This might be due to the fact
 199 that these 100 transformation policies are selected at random out of 127.550 possible options. This is a striking result
 200 nonetheless, which we discuss in-depth in Section 6.

201 **5.2 Results beyond original paper**

202 **Extension 1** We extend the evaluation of the transferability of the searched policies by evaluating the performance
 203 of the policy searched on CIFAR-100 on Rescaled ImageNet. The resulting PSNR values and accuracies are shown
 204 in Table 5. As can be seen from the table, the hybrid strategy produces only 1 dB improvement in PSNR value, and
 205 accuracy decreases by more than 4%. This weakens the claim of transferability made by the authors.

206 **Extension 2** We additionally extend the evaluation of the transferability of the searched policies by testing which
 207 transformations work best on a different dataset. Since good policies share the same general qualities, as stated in Claim
 208 5, the five best transformations from Claim 6 can be expected to be the same when using a different dataset. For this
 209 experiment, we use the Rescaled ImageNet dataset. The resulting transformations are shown in Figure 8c. Out of the 5
 210 best transformations on the Rescaled ImageNet 3 were also found on CIFAR-100 in both our results and the results
 211 from the original paper. This shows that, indeed, these transformations contain the desired qualities from Claim 6.

Policy	PSNR	PSNR (std)	Acc
None	8.96	1.25	61.44
Hybrid	7.92	0.79	57.38

Table 5: PSNR values (dB) and accuracies of policies searched on CIFAR-100 applied to Rescaled ImageNet

Defense	PSNR	PSNR (std)	Acc
Pruning (70%)	11.62	2.18	74.61
Pruning (95%)	10.41	1.32	67.91
Pruning (99%)	9.96	0.57	53.43
Laplacian (10^{-3})	10.73	1.02	71.45
Laplacian (10^{-2})	12.03	0.79	26.20
Gaussian (10^{-3})	12.11	2.98	72.89
Gaussian (10^{-2})	12.13	1.14	36.25

Table 6: Comparisons with existing defense methods under the Adam+Cosine attack

6 Discussion

Overall the results in [5] are reproducible, except Figure 4, with a large discrepancy between our result and the original one - we are still in contact with the authors on this issue. Nevertheless, augmentation policies tend to work as a defense mechanism rather well. For most images, an attacker using reconstruction attacks is unable to find privacy-sensitive information. However, the standard deviation of our results is more than 25% in some settings, and we consider this a valuable metric to contribute. Some images are vulnerable to the attack even with the proposed defense mechanism, and it is as of yet unclear to us which types of images are more vulnerable than others. This issue must be developed further in future research to make the approach widely applicable in real-world use-cases where private data is at stake.

Additionally, we made observations in the codebase that, to the best of our knowledge, were not reported in the paper or any other accompanying documentation. The first was the fact that the loss of the training module was multiplied by a factor of 0.5. This is not a fundamental flaw during the training phase, as it simply produces smaller gradients and therefore leads to a reduced effective learning rate. However, during the reconstruction attacks, the loss used by the attacker was not multiplied by this factor. This makes the attacker in practice use a different loss function from the one used to generate the gradient that it is attempting to match. This may therefore make reconstruction more difficult. Furthermore, we found that two other undocumented augmentations were added in all experiments, namely a random crop and random horizontal flip. Without these, the accuracy of our models decreased by over 10%. We are in contact with the authors regarding these observations, they acknowledged the halved loss as a bug.

6.1 What was easy

The explanation of the general idea and solution of the paper was very clearly put and easy to follow. The codebase contained a README with instructions on how to run some of the paper’s experiments, and these instructions could be followed without significant problems. The code produced results as seen in the paper.

6.2 What was difficult

The most challenging part about reproduction was the unclear description of experiments in the paper and limited clarity in the codebase. Code in the repository was uncommented, used many global variables and many layers of indirection. Many chunks of code were not used, making it harder to follow. Some experimental settings and metrics were not implemented, and some experiment configurations led to fatal errors.

It was very unclear which steps were originally followed to obtain Figure 4. Despite the authors’ helpful comment on which model was used, we were not able to reproduce the correlation, potentially due to randomness in a vast search space (127,550) and the limited sample size (100). Furthermore, the paper does not state how many images were used to produce the PSNR values in the tables. Finally, undocumented augmentations were added in some but not all settings, which was cause for some delay until this was found to be the cause for a 10% accuracy-gap with the authors’ results.

6.3 Communication with original authors

We contacted the authors about multiple clarifications regarding implementation details and notation in the paper. The authors responded promptly and answered almost all of our questions in the first round of contact. We are still in contact on two points. Firstly, regarding our reproduction of Figure 4. Since we got such differing results for this critical part of the authors’ work, we are looking to investigate this further and possibly resolve the discrepancy with them. Secondly, we offered our refactoring of the codebase to the authors as a contribution to their work.

References

- 249
- 250 [1] M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang. Deep learning with
251 differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications*
252 *Security*, pages 308–318, 2016.
- 253 [2] T. S. Brisimi, R. Chen, T. Mela, A. Olshevsky, I. C. Paschalidis, and W. Shi. Federated learning of predictive
254 models from federated electronic health records. *International journal of medical informatics*, 112:59–67, 2018.
- 255 [3] E. D. Cubuk, B. Zoph, D. Mane, V. Vasudevan, and Q. V. Le. Autoaugment: Learning augmentation strategies
256 from data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages
257 113–123, 2019.
- 258 [4] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image
259 Database. In *CVPR09*, 2009.
- 260 [5] W. Gao, S. Guo, T. Zhang, H. Qiu, Y. Wen, and Y. Liu. Privacy-preserving collaborative learning with automatic
261 transformation search. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*,
262 pages 114–123, 2021.
- 263 [6] J. Geiping, H. Bauermeister, H. Dröge, and M. Moeller. Inverting gradients—how easy is it to break privacy in
264 federated learning? *arXiv preprint arXiv:2003.14053*, 2020.
- 265 [7] S. Guo, T. Zhang, X. Xie, L. Ma, T. Xiang, and Y. Liu. Towards byzantine-resilient learning in decentralized
266 systems. *arXiv preprint arXiv:2002.08569*, 2020.
- 267 [8] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE*
268 *conference on computer vision and pattern recognition*, pages 770–778, 2016.
- 269 [9] A. Horé and D. Ziou. Image quality metrics: Psnr vs. ssim. In *2010 20th International Conference on Pattern*
270 *Recognition*, pages 2366–2369, 2010.
- 271 [10] J. Kang, Z. Xiong, D. Niyato, Y. Zou, Y. Zhang, and M. Guizani. Reliable federated learning for mobile networks.
272 *CoRR*, abs/1910.06837, 2019.
- 273 [11] A. Krizhevsky, G. Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- 274 [12] L. Melis, C. Song, E. De Cristofaro, and V. Shmatikov. Exploiting unintended feature leakage in collaborative
275 learning. In *2019 IEEE Symposium on Security and Privacy (SP)*, pages 691–706. IEEE, 2019.
- 276 [13] S. Niknam, H. S. Dhillon, and J. H. Reed. Federated learning for wireless communications: Motivation,
277 opportunities, and challenges. *IEEE Communications Magazine*, 58(6):46–51, 2020.
- 278 [14] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga,
279 A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai,
280 and S. Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural*
281 *Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.
- 282 [15] S. M. Stigler. Francis Galton’s Account of the Invention of Correlation. *Statistical Science*, 4(2):73 – 79, 1989.
- 283 [16] W. Wei, L. Liu, M. Loper, K. H. Chow, M. E. Gursoy, S. Truex, and Y. Wu. A framework for evaluating gradient
284 leakage attacks in federated learning. *CoRR*, abs/2004.10397, 2020.
- 285 [17] H. Xiao, K. Rasul, and R. Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning
286 algorithms, 2017.
- 287 [18] Q. Yang, Y. Liu, T. Chen, and Y. Tong. Federated machine learning: Concept and applications. *ACM Transactions*
288 *on Intelligent Systems and Technology (TIST)*, 10(2):1–19, 2019.
- 289 [19] B. Zhao, K. R. Mopuri, and H. Bilen. idlg: Improved deep leakage from gradients. *arXiv preprint*
290 *arXiv:2001.02610*, 2020.

291 [20] L. Zhu, Z. Liu, and S. Han. Deep leakage from gradients. In H. Wallach, H. Larochelle, A. Beygelzimer,
292 F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32.
293 Curran Associates, Inc., 2019.