
Minimal Model Structure Analysis for Input Reconstruction in Federated Learning

Jia Qian, Hiba Nassar, Lars kai Hansen

Department of Applied Mathematics and Computer Science,
Technical University of Denmark,
2800 Lyngby, Denmark.
{jiaq,hibna,lkai}@dtu.dk

Abstract

Federated Learning (FL) proposed a distributed Machine Learning (ML) framework where every distributed worker owns a complete copy of global model and their own data. The training is occurred locally, which assures no direct transmission of training data. However, the recent work [Zhu et al., 2019] demonstrated that input data from a neural network may be reconstructed only using knowledge of gradients of that network, which completely breached the promise of FL and sabotaged the user privacy.

In this work, we aim to further explore the theoretical limits of reconstruction, speedup and stabilize the reconstruction procedure. We show that a single input may be reconstructed with the analytical form, regardless of network depth using a fully-connected neural network with one hidden node. Then we generalize this result to a gradient averaged over batches of size B . In this case, the full batch can be reconstructed if the number of hidden units exceeds B . For a Convolutional Neural Network (CNN), the number of required kernels in convolutional layers is decided by multiple factors, e.g., padding, kernel and stride size, etc. We require the number of kernels $h \geq (\frac{d}{d'})^2 C$, where we define d as input width, d' as output width after convolutional layer, and C as channel number of input. We validate our observation and demonstrate the improvements using bio-medical (fMRI, White Blood Cell (WBC)) and benchmark data (MNIST, Kuzushiji-MNIST, CIFAR100, ImageNet and face images).

Federated Learning [Konečný et al., 2016] was proposed by Google for distributed network training and has created tremendous interest. One important virtue of FL is to keep data on the generating-device to avoid information leakage due to direct data transmission. However, the distributed setup opens to exposure for attackers, and sensitive information may be disclosed in other forms. Recent work [Zhu et al., 2019] demonstrated a severe attack - reconstruction of input data within an FL environment. This idea was similar to model inversion [Fredrikson et al., 2015]; however, it is easy and effective to apply this attack in a federated setup. By imitating an honest server, data may be reconstructed with the access to gradient updates and model parameters. It seriously breaches the promise of FL - data is kept locally. Here we aim to explore the limits and efficiency of this attack. We offer theoretical analysis of the reconstruction based on both a fully-connected neural network and CNN, from a perspective of solving a linear system. Our main contributions are summarized as the follows.

- We show that reconstruction of input amounts to solving a set of linear equations, and based on which we give minimal structural conditions for reconstruction based on the fully-connected neural network (a.k.a. Multilayer Perceptron (MLP)) and CNN.
- We show that MLP only needs a **single** node in **one** hidden layer to reconstruct a single input image, regardless of the model depth. Complementing Zhu et al. [2019], we derive a

closed form for a lossless reconstruction in this case. Moreover, we generalize the result to batch reconstruction and show that for a full reconstruction the number of hidden units has to exceed the batch size.

- We show that for CNN, the number of kernels in every convolutional layer should be such that the size of output after passing through the convolutional layers exceeds the size of original input.
- We propose a reconstruction method in response to three cases: 1) single-instance reconstruction using MLP; 2) single-instance reconstruction using CNN; 3) batch reconstruction.
- We also suggest to include a regularizer for batch reconstruction, which increases the numerical stability during iterative optimization and outputs more faithful reconstructions.

The paper is organized as follows: in Section 1, we briefly introduce FL and present key related works and in Section 2 we focus on our reconstruction method and offer the theoretical analysis based on the model architecture. Finally, we provide numerical demonstrations in Section 3 and conclude in Section 4.

1 Background and related works

1.1 Federated Learning

FL is a distributed ML paradigm that incorporates a set of distributed workers (customers) and a server to jointly train a global model. Unlike traditional Cloud-based ML [Yang et al., 2019], it has no (training) data transmission between workers and the server. Instead, the set of workers collaborate to optimize the global cost function that is estimated by a collection of local cost functions. The workers own their data, the full copy of the global model, and implement local training (minimizing local cost functions) using their data and update the gradients with the server. Without loss of generality, we assume that each worker i owns data generated from the empirical distribution \hat{D}_i , and the data owned by different workers might be heterogeneous. Given p active workers participate in each round, and m_i is the amount of data on worker i and m is the sum of all data $m = \sum_{i=1}^p m_i$. All the workers share the same batch size B . The global empirical loss function $\hat{\ell}$ can be approximated by the weighted combination of local loss functions $\hat{\ell}_i$. It is defined as:

$$\hat{\ell}(x, y) = \sum_{i=1}^p \frac{m_i}{m} \mathbb{E}_{(x, y) \sim \hat{D}_i} [\hat{\ell}_i(f_w(x), y)] \quad (1)$$

Each iteration consists of two stages: local training (workers) and aggregation (server). More specifically, after distributed workers have completed local training, they share the corresponding gradients with the server, and the server aggregates the gradients based on the given criterion and finally it sends the updated global model (or aggregated gradients) back to the workers. The workers will use the updated model for the next iteration. This procedure can be repeated for multiple times.

1.2 Related work

The feasibility of input reconstruction based on the gradients and parameters of the neural works was demonstrated by Zhu et al. [2019]. They proposed to reconstruct input from an initial guess sampled from normal distribution and iteratively optimize the reconstruction by minimizing the distance between gradients and guessed gradients. However, it does not give strong experimental results on larger batches. Zhao et al. [2020] expand on Zhu et al. [2019] to improve the label prediction accuracy. Geiping et al. [2020] proposed a new cost function and reconstruct the input based on deep neural networks like ResNet [He et al., 2016]. Wei et al. [2020] introduced a framework for evaluation of privacy leakage and relation to FL hyperparameters. Pan et al. [2020] also focus on deep neural network like VGG [Simonyan and Zisserman, 2014], GoogLeNet [Szegedy et al., 2015], which are normally are more informative for the reconstruction. We are interested in minimal requirements analysis of reconstruction based on the network structure which is currently under-explored.

2 Reconstruction method and theoretical analysis

In this section, we will first present the reconstruction method and then introduce the theoretical analysis of minimal reconstruction conditions based on the two architectures: MLP and CNN.

2.1 Reconstruction method

Input reconstruction is essentially an inverse problem, without the loss of generality, say we have any function G parameterized by w , which maps from input $x \in R^d$ to a (gradient) vector v , thus $G_w : R^d \rightarrow v$. For the batch case, the expected gradient is $\bar{v} = \frac{1}{B} \sum_{i=1}^B G(x_i, y_i; w)$. We aim to compute x with the knowledge of model parameters w and output v (or \bar{v}). If G^{-1} exists, we can compute x directly, which is typically not the case. The intriguing question is that is it still possible to reconstruct x in the noninvertible case? We show that in Section 2.2, one-instance reconstruction based on MLP has an analytical form and can be computed directly. This is different from Zhu et al. [2019] where they consider it as noninvertible and address it as an optimization problem. For one-instance CNN reconstruction, we propose a two-step method where we first compute the output of convolutional layer using the closed-form mentioned before, and then we apply the optimization method. While, for batch reconstruction, there is no analytical form in most cases. We convert it to the optimization problem; more specifically, we start from a random guess (\hat{x}, \hat{y}) , and gradually pull it to close to the original input x by minimizing the cost function, like proposed by Zhu et al. [2019]. For the batch reconstruction, we augment the cost function by an additive regularizer as $L(\cdot) + \lambda R(\cdot)$ where $L(\cdot)$ is the distance between ground-truth gradients and guessed gradients. For $R(\cdot)$ we suggest an orthogonality regularizer defined as $R = \lambda \sum_{k \neq k'=1}^n (\hat{x}_k^T \hat{x}_{k'})^2$, if \hat{x}_k and $\hat{x}_{k'}$ are orthogonal the product is zero, which implies that the optimizer promotes solutions where the batch members are dissimilar. Or a L2 regularizer defined as $\lambda \sum_i \hat{x}_i^T \hat{x}$. We experimentally explore the critical role of the regularizer for large batch size in Section 3.

Algorithm 1 Insecure FL with reconstruction attack

```

1: Initialization:  $w^0$ 
2: for  $t=1, \dots, T$  do
3:   => workers:
4:   for  $j=1, 2, \dots, p$  do  $p$  workers (in Parallel)
5:      $v_j^t = \nabla \hat{\ell}_j(f(X_j^t; w^t), Y_j^t)$ 
6:     with  $(X_j^t = \{x_{jk}^t\}_{k=1, \dots, B}, Y_j^t = \{y_{jk}^t\}_{k=1, \dots, B}) \sim \hat{D}_j$ 
7:     share  $v_j^t$  with server
8:   end
9:   => server(attacker):
10:   $w^{t+1} = w^t - \eta \times \frac{m_j}{m} \sum_{j=1}^p v_j^t$ 
11:  share  $w^{t+1}$  with workers for next round
12:   $\hat{X}_j^t, \hat{Y}_j^t = \text{Reconstruction}(v_j^t, w^t, m, \lambda, \text{prior})(\text{Algorithm 2})$ 
13: end

```

We show insecure FL with potential reconstruction attack in Algorithm 1 where the attacker could be the server who has the complete knowledge of gradients update and model parameters from each worker j at each round t . In Algorithm 2, we present the reconstruction approaches in response to three cases: 1) one-instance MLP reconstruction, 2) one-instance CNN reconstruction, and 3) batch reconstruction (using MLP and CNN). The pseudo code in Algorithm 3 demonstrates the iterative optimization step.

Our method implementation deviates from the pioneering work [Zhu et al., 2019] in a few ways. First, we derive a closed-form for one-instance MLP reconstruction, which leads to a faster and more accurate reconstruction (almost lossless). Second, we divide one-instance CNN reconstruction into two steps; first, we directly compute the output of convolutional layer using the advantage of closed-form and based on which we reconstruct the input (a.k.a deconvolution), which speeds up the overall reconstruction. Last, we expand the cost function with either an orthogonality regularizer or L2 regularizer for batch reconstruction. The orthogonality regularizer may penalize the similarities between reconstructed images (since we only know the average gradient of the batch), mainly when

Algorithm 2 Reconstruction

```
1: Input:  $v, w, m, \lambda$ , prior
2: if 1) single recon. & mlp then ▷ case one
3:   Return  $\hat{x} = \frac{\partial \ell}{\partial w_{1i}^1} / \frac{\partial \ell}{\partial b_1^1} \quad \forall i < d$  ▷  $w_{1i}^1, b_1^1$  are the weights and bias in 1st hidden layer,  $d$  is input dimension
4: End
5: if 2) single recon. & cnn then ▷ case two
6:    $\hat{z} = \frac{\partial \ell}{\partial w_{1j}^1} / \frac{\partial \ell}{\partial b_1^1} \quad \forall j < d'$  ▷  $w_{1j}^1, b_1^1$  are the weights and bias on 1st hidden layer after conv. layer,  $d'$  is dimension of output of conv. layer
7:   Return  $\text{ltr\_rec}(\hat{z}, w_{\text{partial}}, \text{"partial"}, m, \lambda, \text{prior})$  ▷  $w_{\text{partial}}$  refers to the params. of convol. layer
8: End
9: if 3) batch reconstruction then ▷ case three
10: Return  $\text{ltr\_rec}(v, w, \text{"all"}, m, \lambda, \text{prior})$ 
11: End
```

image patterns are similar (e.g., MNIST). In contrast, L2 regularizer offers faster convergence, in particular, when batch size is large. Our main focus in this work is to explore and validate the minimal condition of the full reconstruction.

2.2 Reconstruction with fully-connected neural network

Algorithm 3 ltr_rec

```
1: Input:  $v, w$ , flag,  $m, \lambda$ , prior
2: for  $i=0,1,2,\dots,I$  do  $I$  iterations
3:   if  $i==0$  then
4:     if prior==uniform then
5:        $\hat{X}_0, \hat{Y}_0 \sim \mathcal{U}(0, 1)$ 
6:     if prior==normal then
7:        $\hat{X}_0, \hat{Y}_0 \sim \mathcal{N}(0, 1)$ 
8:      $\hat{v}_1 = G(\hat{X}_0, \hat{Y}_0; w)$  ▷  $G$  is function of computing gradients
9:     else
10:       $\hat{v}_i = G(\hat{X}_i, \hat{Y}_i; w)$ 
11:      if flag=="all" then
12:         $L = \|v - \hat{v}_i\|_2^2 + \lambda R(\hat{X}_i)$  ▷  $R$  is the regularizer
13:      else
14:         $L = \|v - \hat{v}_i\|_2^2$ 
15:      end
16:      update:  $\hat{X}_{i+1} = \hat{X}_i - \eta \times \nabla_{\hat{X}_i} L$ 
17:              $\hat{Y}_{i+1} = \hat{Y}_i - \eta \times \nabla_{\hat{Y}_i} L$ 
18:      if  $(i/m)==0$  then
19:         $\lambda = 0.9 * \lambda$ 
20:      end
21:    end
22:  return  $\hat{X}_I, \hat{Y}_I$ 
```

Say we have a one-layer MLP f , with n_1 units in hidden layer and n_2 units in output layer (if classification task, n_2 is equal to the number of classes). Thus, we can define the output of model as $a_j = f_w(x) = \sum_{i=1}^{n_1} w_{ji}^2 \sigma(w_i^1 x + b_i^1) + b_j^2 \quad \forall j \in [1, n_2]$, where w^1 and b^1 are the weights and bias in hidden layer, and w^2, b^2 in output layer and $\sigma(x)$ is sigmoid (monotonic) activation function. For the classification task, we employ cross-entropy as the cost function $\ell(p_i, y_i) = -\sum_j^C y_{ij} \log p_{ij}$ where p_{ij} is the output of softmax function, and y_i is the one-hot encoding vector with all zeros except the corresponding class indicating one.

Proposition 1 (ONE-INSTANCE MLP RECONSTRUCTION). *To reconstruct one input based on MLP, we derive the analytical form to compute the (almost) lossless input, with only **single** unit in the first hidden layer as long as bias term exists, regardless how deep the network is.*

The comprehensive proof can be found in appendix 6. It can be generalized to deep fully-connected neural network (say last layer is L). We only need two ingredients in our recipe, the partial derivative w.r.t the bias term and weights in the first hidden layer (after input layer). It shows that training with Stochastic Gradient De-

scendent (SGD) is very vulnerable to reconstruction attack in FL (distributed ML), in particular, on fully-connected neural network no matter how deep it is.

Proposition 2 (BATCH MLP RECONSTRUCTION). *For batch reconstruction, the number of units in first hidden layer should meet $n_1 \geq B$, given the high-dimension input whose dimension $d \gg n_2, d \gg B$.*

Proof. (sketch:) Reconstructing $x_1, x_2, \dots, x_B \in \mathbb{R}^d$ approximates to solving the linear equations. Given an invertible sigmoid function and a single output y_i , we may compute the unique $\sigma'(x_i)$. The number of equations exceeds the number of variables $n_2 + n_1 n_2 + n_1 + n_1 d \geq B n_2 + n_1 B + B d$, thus $n_1 \geq \frac{B n_2 + B d - n_2}{n_2 + 1 + d - B}$. Typically we have $d \gg n_2, d \gg B$, and it is dominated by $\frac{B(d+n_2)}{d+n_2}$, therefore $n_1 \geq B$. \square

In general, batch reconstruction is more challenging since we typically only know the average or the sum of the gradients and aim to reconstruct every individual instance. We refer to this procedure as *demixing* in the following context. Theoretically, if the number of equations is identical or greater than batch size and all the equations are independent, it is solvable. However, it is not easy to solve in practice, mainly when x is high-dimensional and the scales in the linear system are small. We use an iterative method to solve it. Besides, the to-be-optimized input value during optimization might introduce the saturation of sigmoid function, i.e., $\sigma(x) \approx 0, \sigma(x) \approx 1$ regardless the change of x outside the interval $[-4, 4]$. The choice of optimization method is important. For a high-dimensional (non-sparse) input, second-order or quasi-second-order methods sometimes fail since they are designed to search for the zero-gradient area. The number of saddle points exponentially increases with the number of input dimensions [Dauphin et al., 2014]. The first-order method, e.g., Adam [Kingma and Ba, 2014] takes much more iterations to converge, but it is relatively more stable, likely to escape from the saddle points [Goodfellow et al., 2016].

2.3 Reconstruction with convolutional neural network

Proposition 3 (SINGLE-LAYER CNN RECONSTRUCTION). *To reconstruct a single-layer CNN immediately stacked by a fully-connected layer, $h \geq (\frac{d}{d'})^2 C$ kernels are required, where C is the channel number of input, d is the width of input, and d' is width after convolutional layers.*

For a single convolutional layer CNN, the kernel parameters (kernel size k , padding size p , stride size s) determine the output size d' after convolutional layers. Say we have input $X \in \mathbb{R}^{B \times C \times d \times d}$, for the simplicity we assume height and width are identical, and C is the channel number and B is the batch size. We define a square kernel with width k (weights indicated as l^0), bias term r , and we have h kernels. After convolutional layer, the width of output is $d' = \frac{d+2p-k}{s} + 1$. The output of convolutional layer is defined in eq. (2).

$$z_{mij} = \left(\sum_{c=1}^C \sum_{g=1}^k \sum_{n=1}^k l_{mcgn}^0 \hat{x}_{c, si+g-1, sj+n-1} \right) + r_m \quad (2)$$

$$\forall (i, j) \in [1, d'] \times [1, d'], \forall m \in [1, h]$$

We include the proof in appendix 7. The number of equations should be equal or greater than the number of unknowns. From eq. (2), we have $(d')^2 B h$ equations and $d^2 B C$ unknowns, thus we need $h \geq (\frac{d}{d'})^2 C$, with the assumption that H is known, which can be solved by enough units in dense layer from Proposition 2. One special case is that convolutional layer is stacked by an output layer directly (no dense layer), then we need to meet $h \geq (\frac{d}{d'})^2 C B$, $n_0 \geq \frac{n_1(B-1)}{n_1-B}$, and $1 < B < n_1$ to solve H since σ is monotonic activation function ($n_0 = h(d')^2$). Note here n_1 indicates the number of units in output layer (as no dense layer). It can also be generalized to multiple convolutional layers, more details can be found in appendix 7.

CNN reconstruction can be seen as a two-stage reconstruction: *demixing* and *deconvolution*. Namely, the demixing stage is essentially an inverse procedure of fully-connected network and we aim to reconstruct the output of convolutional layer (input of fully-connected layer). The deconvolution stage we want to reverse the convolutional step, starting from the output of convolutional layer to reconstruct the original input. For the demixing stage, we can apply the conclusion from MLP, as long as the number of hidden units is equal or greater than the batch size, then theoretically we may evaluate the individual instance.

3 Experimental results

Dataset and Setup. MNIST and KMNIST are the handwritten digits and Kuzushiji accordingly, with size $1 \times 28 \times 28$ and it totally contains 10 classes. CIFAR100 contains 100 classes and every class has 600 images with size $3 \times 32 \times 32$. ImageNet contains 1000 classes, and each image has size $3 \times 64 \times 64$. Every image in WBC dataset has size $3 \times 240 \times 320$ and four classes (Eosinophi,

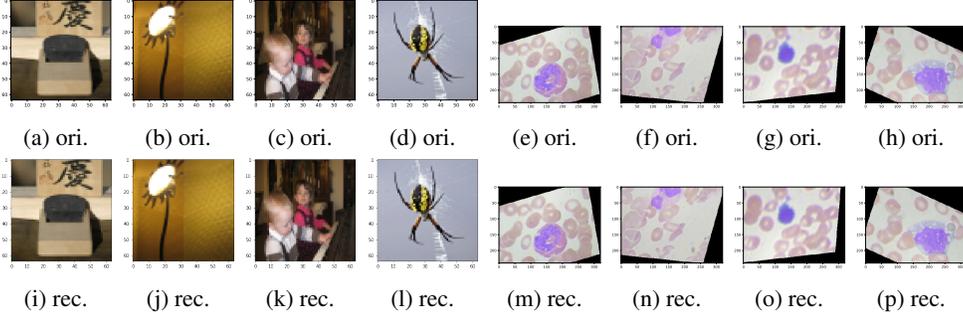


Figure 1: ImageNet and White Blood Cell analytical reconstruction. The first row includes the original inputs and the second row corresponds to all the reconstructions.

Neutrophil, Lymphocyte and Monocyte). We have implemented our method in Python and run it on a Titan X GPU with Architecture Maxwell and 12 GB of Ram.

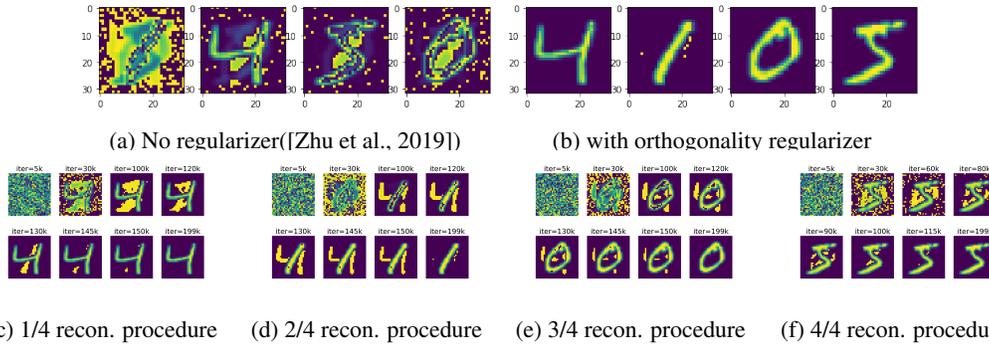


Figure 2: *MNIST*: 2a and 2b show the final reconstruction without and with regularizer accordingly. Moreover, the reconstructions procedures of 2b are separately shown in 2c, 2d, 2e and 2f.

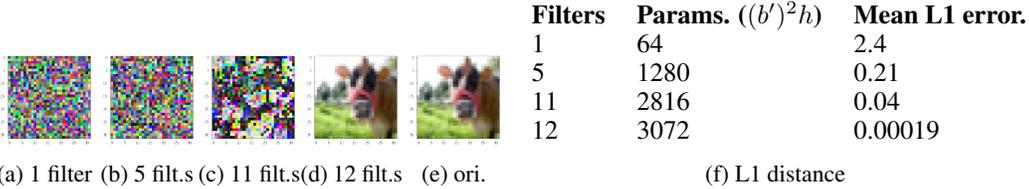


Figure 3: *One-layer CNN*: we show that the image reconstruction changes with increasing number of filters (numerical result is shown in Table 3f).

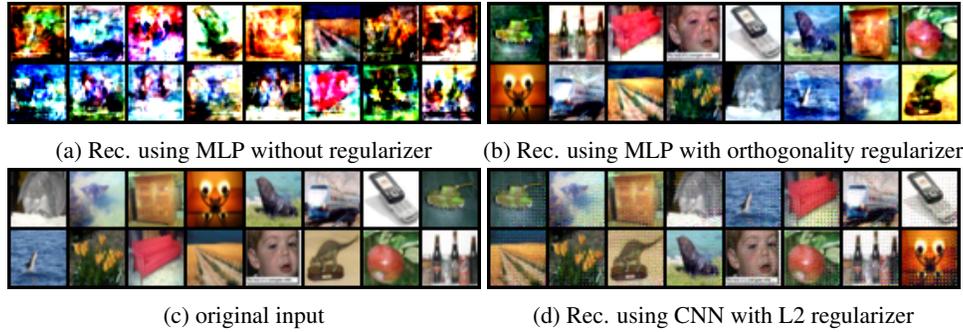


Figure 4: Batch size 16 (MLP and CNN).

MLP Reconstruction Analysis. We implement one-instance reconstruction on MLP by the analytical form $x_i = \frac{\partial \ell}{\partial w_{1i}} / \frac{\partial \ell}{\partial b_1^1}$, $\forall i \in [1, d]$, using one-hidden layer MLP with only one unit in hidden layer. The outcome is demonstrated in Figure 1 using ImageNet and WBC dataset, the average L1 distance

per pixel between original input and reconstruction is bounded below $1e - 8$, which is almost lossless with $\mathcal{O}(n)$ complexity. For MLP batch reconstruction, we experimentally set the batch size equal to four, identical to the number of units in hidden layer, as shown in Figure 2. We first give the final reconstruction of four inputs without regularizer in Figure 2a, whereas in Figure 2b the reconstruction quality is being improved significantly with the orthogonality regularizer. More specifically, we start λ from 0.1 and gradually decay after 200 epochs by 90%. Besides, from Figure 2c to 2f, we partially show the reconstruction procedure and we can see how the regularizer plays the key role during optimization procedure to hinder the similarities between instances. Moreover, we also show MLP reconstruction with batch size 16 in Figure 4a and 4b. Empirically, it shows that the optimization for high-dimension demixing is very challenging and the existence of regularizer significantly improves the numerical stability and the reconstruction quality. In Figure 5, we demonstrate that Proposition 2 is also valid when batch size B is large, eg., $B = 100$.

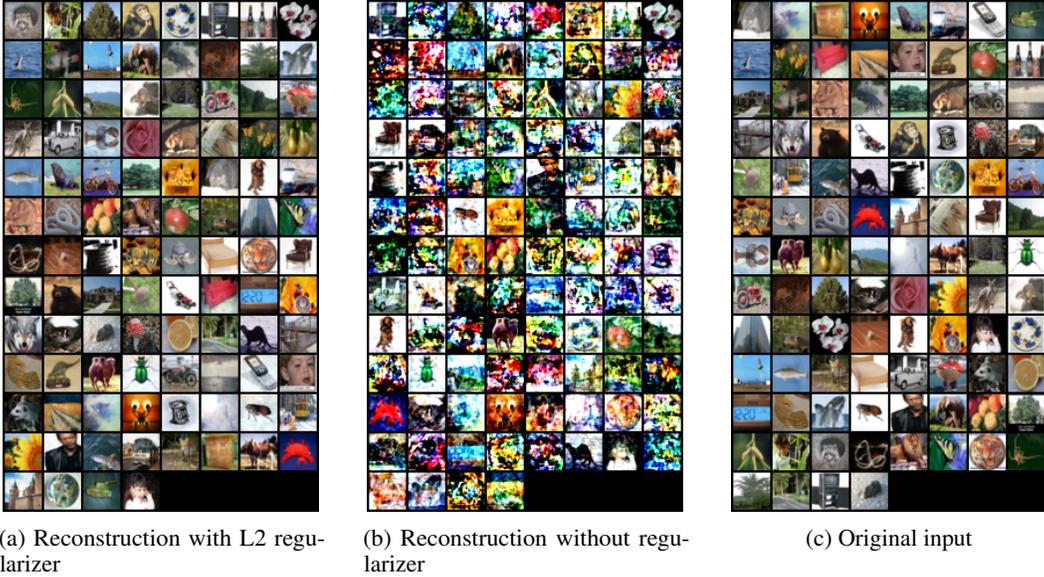


Figure 5: MLP reconstruction with batch size 100 (CIFAR100).

CNN Reconstruction Analysis. We divide one-instance CNN reconstruction into two steps: 1) we directly compute the output of convolutional layer; 2) apply iterative optimization. We test the reconstruction performance with different numbers of kernels. We let the kernel size be 5, padding size 2, stride size 2 and pooling size 1, thus at least 12 kernels are required according to Proposition 3. In Figure 3 we first visually show the reconstruction improvement with the incremental numbers of kernels, and from Table 3f we numerically show reconstruction error (L1 distance) decreasing with more kernels. For batch CNN reconstruction, we demonstrate it in Figure 4c and 4d where we have the similar convolutional setup with one-instance CNN, thus 12 kernels are applied in convolutional layer. Note for the redundant architectures, i.e., with more kernels and units than required, we can easily mask the corresponding gradients during iterative optimization to fasten the reconstruction step.

4 Conclusion

We theoretically studied the minimal structural requirements for reconstruction and analyzed the relations between network architecture, size, and reconstruction quality. We show that the number of units in first hidden layer should be equal to greater than batch size using MLP. It is worthwhile to mention that joint training using MLP (with bias) with batch size equal to one in FL is extremely vulnerable to the adversary. For CNN, the number of kernels along with the number of units in a fully-connected (dense) layer decides the quality of reconstruction. Specifically, number of units in hidden layer is determined by batch size and required number of kernels is decided by output dimension after convolutional layers. Our observations also apply to big batch size with the aid of the regularizer. We hope that the limits explored in the present work and conditions for reconstruction can aid the practitioners to choose network architecture and communication strategies when applying FL on sensitive information-related applications, e.g., medical diagnosis, stocking price prediction.

5 Acknowledgement

The research leading to these results has received funding from the European Union’s Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement No. 764785, FORA-Fog Computing for Robotics and Industrial Automation.

References

- Ligeng Zhu, Zhijian Liu, and Song Han. Deep leakage from gradients. In *Advances in Neural Information Processing Systems*, pages 14774–14784, 2019.
- Jakub Konečný, H Brendan McMahan, Felix X Yu, Peter Richtárik, Ananda Theertha Suresh, and Dave Bacon. Federated learning: Strategies for improving communication efficiency. *arXiv preprint arXiv:1610.05492*, 2016.
- Matt Fredrikson, Somesh Jha, and Thomas Ristenpart. Model inversion attacks that exploit confidence information and basic countermeasures. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, pages 1322–1333, 2015.
- Qiang Yang, Yang Liu, Tianjian Chen, and Yongxin Tong. Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 10(2):1–19, 2019.
- Bo Zhao, Konda Reddy Mopuri, and Hakan Bilen. idlg: Improved deep leakage from gradients. *arXiv preprint arXiv:2001.02610*, 2020.
- Jonas Geiping, Hartmut Bauermeister, Hannah Dröge, and Michael Moeller. Inverting gradients—how easy is it to break privacy in federated learning? *arXiv preprint arXiv:2003.14053*, 2020.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- Wenqi Wei, Ling Liu, Margaret Loper, Ka-Ho Chow, Mehmet Emre GURSOY, Stacey Truex, and Yanzhao Wu. A framework for evaluating gradient leakage attacks in federated learning. *arXiv preprint arXiv:2004.10397*, 2020.
- Xudong Pan, Mi Zhang, Yifan Yan, Jiaming Zhu, and Min Yang. Theory-oriented deep leakage from gradients via linear equation solver. *arXiv preprint arXiv:2010.13356*, 2020.
- Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9, 2015.
- Yann N Dauphin, Razvan Pascanu, Caglar Gulcehre, Kyunghyun Cho, Surya Ganguli, and Yoshua Bengio. Identifying and attacking the saddle point problem in high-dimensional non-convex optimization. In *Advances in neural information processing systems*, pages 2933–2941, 2014.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Ian Goodfellow, Yoshua Bengio, Aaron Courville, and Yoshua Bengio. *Deep learning*, volume 1. MIT Press, 2016.