Federated Learning for Biometrics Applications

Vishal M. Patel
Associate Professor
Department of Electrical and Computer Engineering
JOHNS HOPKINS UNIVERSITY

vpatel36@jhu.edu
https://engineering.jhu.edu/vpatel36/

January 12, 2023

Agenda

Part 1

- Motivation
- Federated learning
  - FedAvg
  - SplitNN
- Privacy-enhancing methods for federated learning

Part 2

- Applications
  - Face recognition
  - Face anti-spoofing
  - Active authentication
  - Thermal to visible face synthesis
- Open problems

ImageNet Challenge

- Large Scale Visual Recognition Challenge (ILSVRC) 2017
  - 1000 object categories
  - 1.2M training images

Classification Results (CLS)

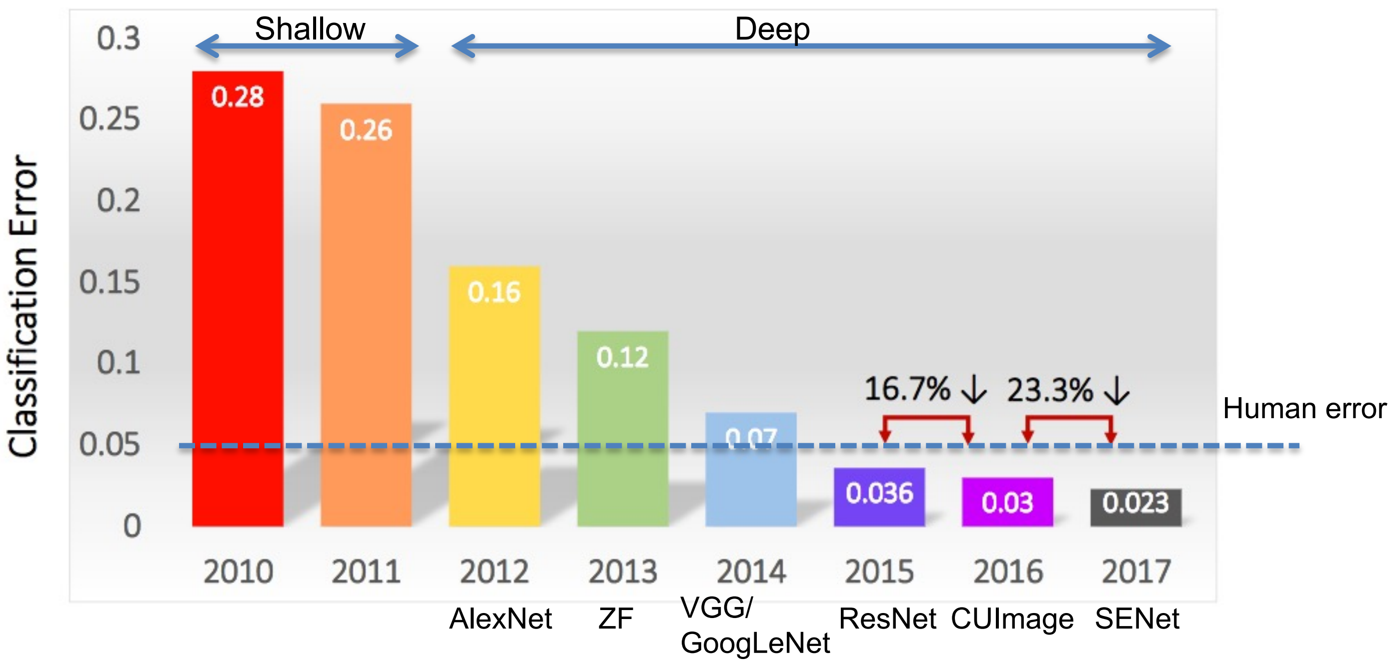

Russakovsky et al. ILSVRC2017 CVPRW 2017

Face Recognition

- Labeled Faces in the Wild (LFW)
  - 5,749 subjects
  - 13,233 faces

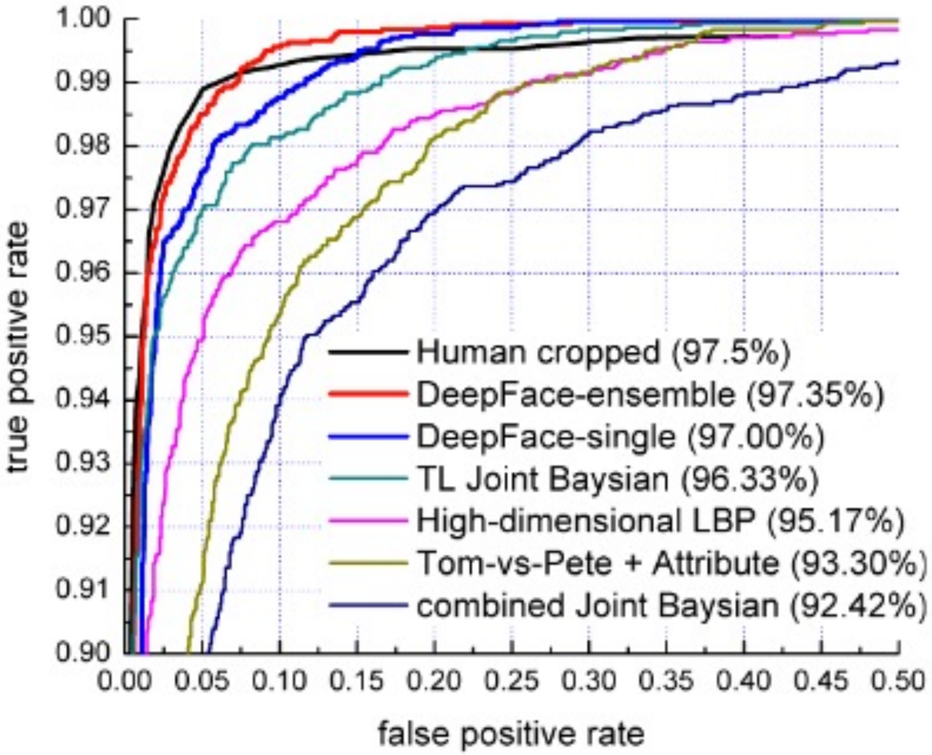

- Mean classification accuracies:

  YI+AI (0.9983 ± 0.0024)

  FRDC (0.9972 ± 0.0029)

  CHTFace (0.9960 ± 0.0025)

Training data: 4 million faces, 4000 identities (facebook)

Taigman et al. DeepFace: Closing the Gap to Human-Level Performance in Face Verification, CVPR 2014 .

Detectron – Facebook

- Detectron model for object detection
  - Trained on a large-scale image data from Instagram

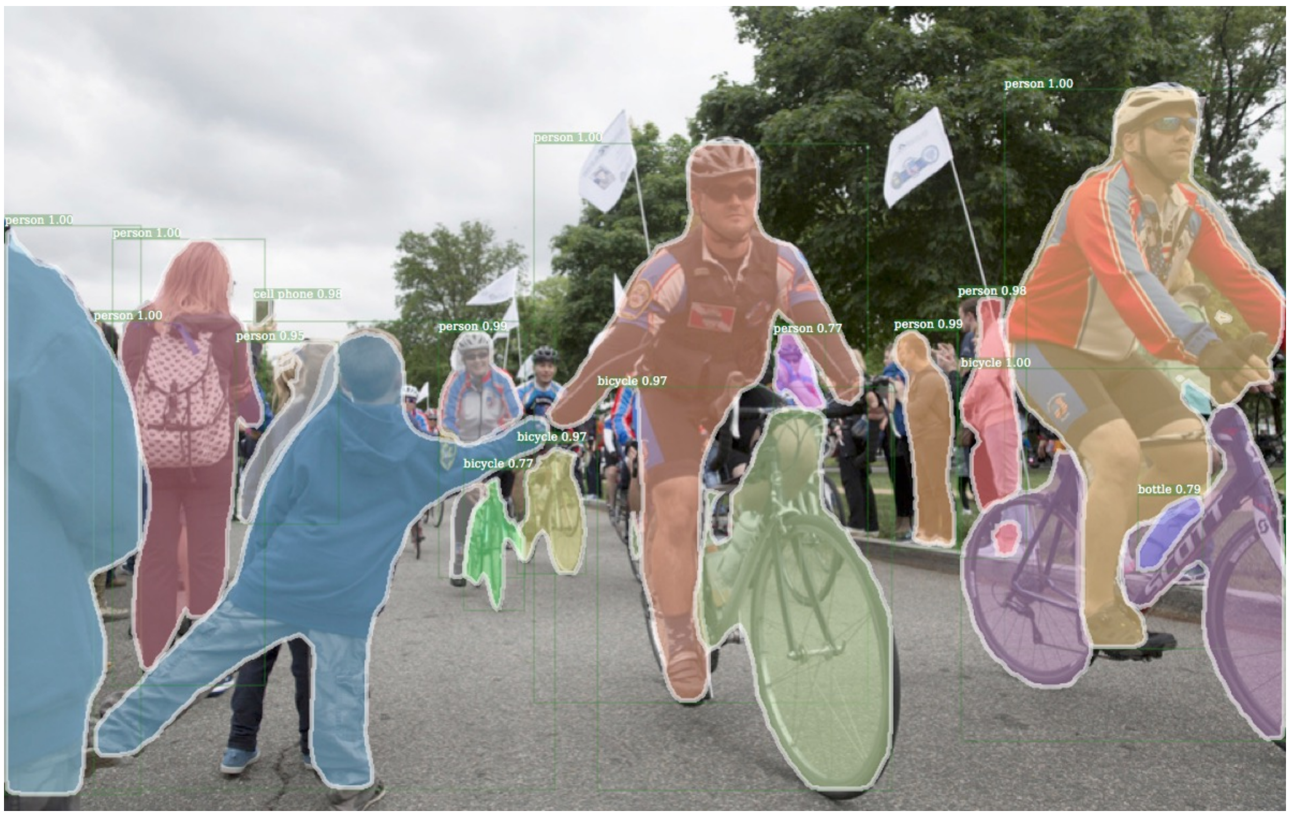

https://github.com/facebookresearch/detectron

LeNet5 vs AlexNet

LeNet5 LeCun et al. 1998

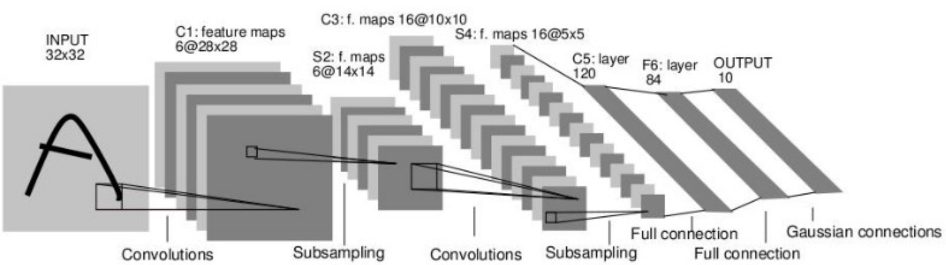

- Trained on MNIST digit dataset with 60K training examples
- Sigmoid or tanh nonlinearity
- Average pooling
- Fully connected layers at the end

AlexNet Krizhevsky et al. 2012

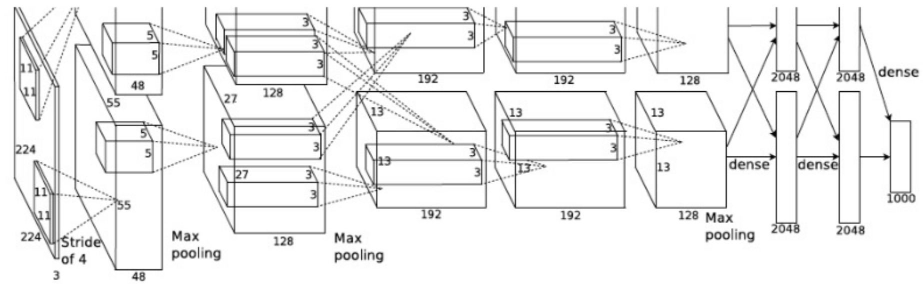

- Trained on ImageNet dataset with 1.2M training images
- Rectified Linear Unit (ReLU) nonlinearity
- Max pooling
- GPU implementation
  - Trained on two GPUs for a week
- Dropout regularization
- Fully connected layers at the end

Why?

- Availability of large annotated data
- More layers
  - Capture more invariances
- More computing
  - Availability and affordability of GPUs
- Better regularization
  - Dropout
- New nonlinearities
  - Rectified Linear Unit (ReLU)
  - Parametric Rectified Linear Unit (PReLU)

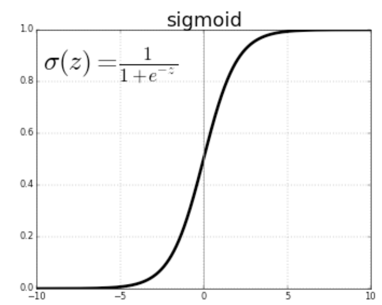

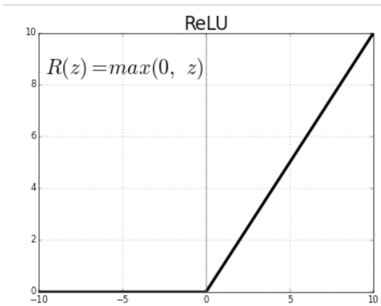

Razavian et al. CVPR 2014

Why?

- Availability of large annotated data
- More layers
  - Capture more invariances
- More computing
  - Availability and affordability of GPUs
- Better regularization
  - Dropout
- New nonlinearities
  - Rectified Linear Unit (ReLU)
  - Parametric Rectified Linear Unit (PReLU)

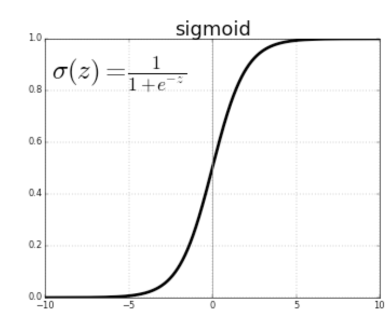

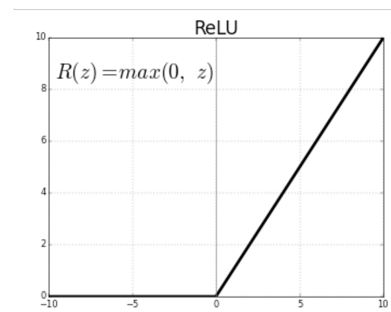

Razavian et al. CVPR 2014

Large Datasets

- Collecting and annotating datasets
  - Expensive
  - Labor intensive
  - User privacy issues
    - GDPR: General Data Protection Regulation
    - HIPAA: Health Insurance Portability and Accountability Act, 1996
    - SHIELD: Stop Hacks and Improve Electronic Data Security Act, Jan 1 2019
    - PCI: Payment Card Industry Data Security Standard, 2004
    - IRB: Institutional Review Board

Protecting User Privacy

- Data privacy (protect the data)
  - Cancelable biometrics
    - Modify data through revocable and non-invertible transformations
  - BioHashing
    - Random projections are used to generate templates
  - Differential privacy
    - An algorithm is differentially private if its behavior hardly changes when a single individual joins or leaves the dataset
    - Hide unique samples (add noise to data)
  - Homomorphic encryption
    - Perform calculations on encrypted data
- Federated learning (build protection into the models)
  - Machine learning on decentralized data
  - *Communication-efficient learning of deep networks from decentralized data*, AISTATS 2017, McMahan et al. (Google)

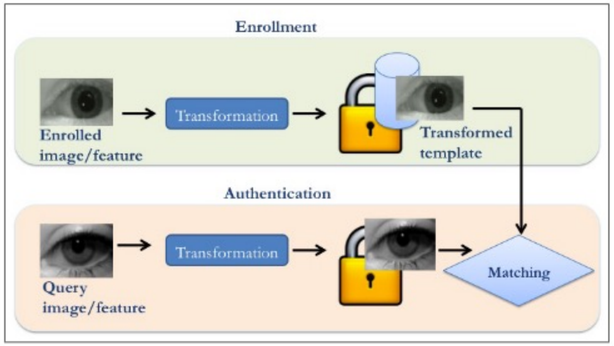

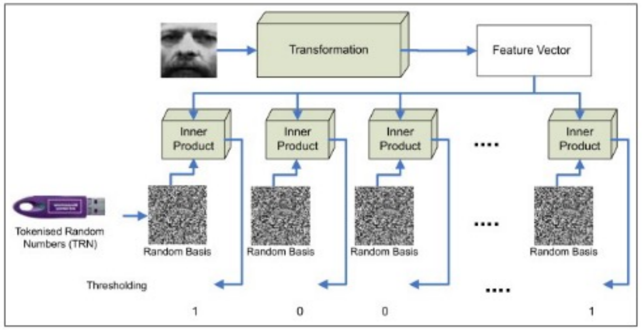

V. M. Patel, N. K. Ratha and R. Chellappa, "Cancelable Biometrics: A review," in IEEE Signal Processing Magazine, vol. 32, no. 5, pp. 54-65, Sept. 2015.

Federated Learning - FedAvg

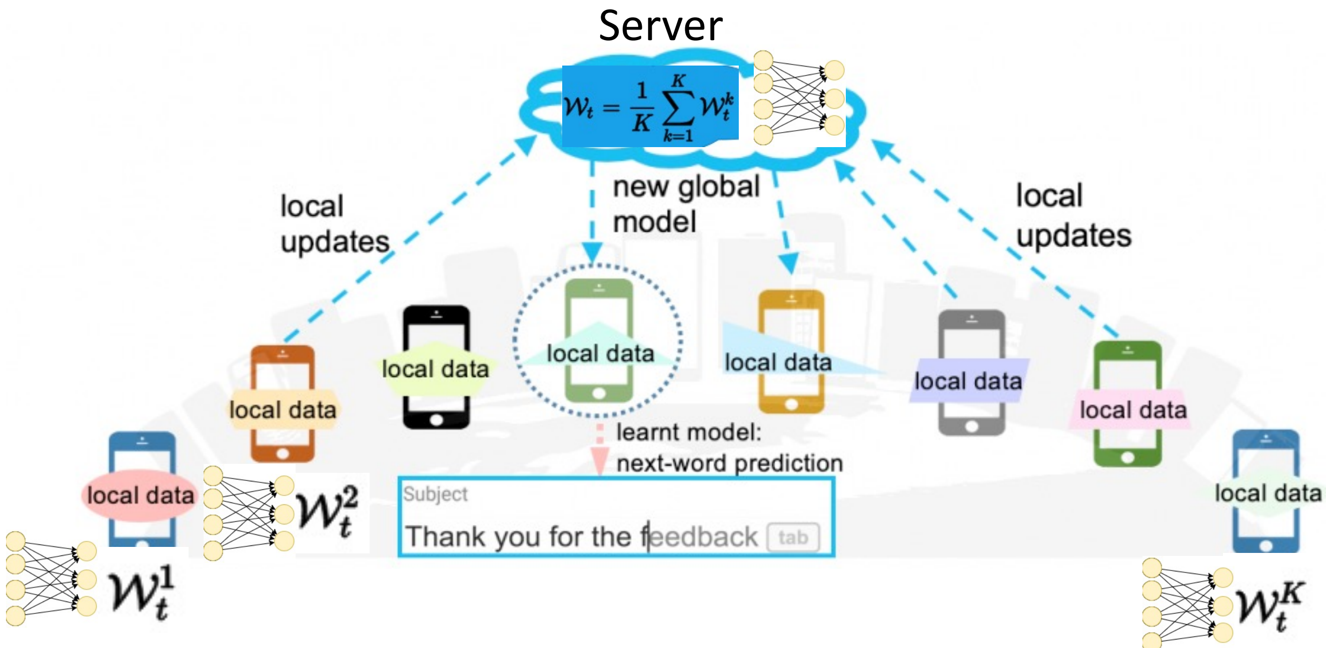

- Different users (clients) collaboratively learn a machine learning model with the help of a server
- Local training
  - Users locally compute training parameters and send them to the server
- Model aggregating
  - The server performs secure aggregation over the uploaded parameters from different users without learning local information
- Parameters broadcasting
  - The server broadcasts the aggregated parameters to the users
- Model updating
  - All users update their respective models with aggregated parameters and test the performance of the updated models

Li et al. IEEE SPM 2021

Federated Learning - Applications

- Learning over smart phones
  - Mobile-based biometrics applications
  - Active authentication
- Learning across organizations
  - Multi-institutional collaboration
- Internet of things
  - Wearable devices, autonomous vehicles, smart homes, …

Federated Learning - Applications

- Next word prediction (Google)
  - Federated Learning for Mobile Keyboard Prediction, Hard et al., 2018
- Speaker recognition (Apple Siri)
  - QuickType (Apple's personalized keyboard)

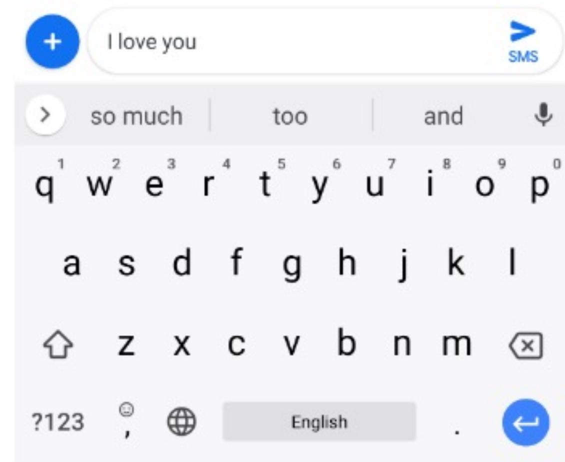

Fig. 1. Next word predictions in Gboard. Based on the context "I love you", the keyboard predicts "and", "too", and "so much".

MIT Technology Review

Artificial intelligence / Machine learning

How Apple personalizes Siri without hoovering up your data

The tech giant is using privacy-preserving machine learning to improve its voice assistant while keeping your data on your phone.

by **Karen Hao** December 11, 2019

Federated Learning - Challenges

- Communication
  - Federated networks are comprised of a massive number of devices which causes communication in the network to be slower than local computations (i.e. expensive communication)
  - Need communication-efficient methods that iteratively send model updates as part of the training process
- Systems heterogeneity
  - Storage, computational, and communication capabilities of each device in federated networks may differ due to variability in hardware (CPU, memory), network connectivity (3G, 4G, 5G, wifi), and power (battery level)
  - Stragglers and fault tolerance significantly more prevalent
- Non-IID data
  - Devices frequently generate and collect data in a non-identically distributed manner across the network.
  - Unbalanced data
  - Increases the likelihood of stragglers, and may add complexity in terms of modeling, analysis, and evaluation
- Privacy issues

Federated Learning – Privacy Issues

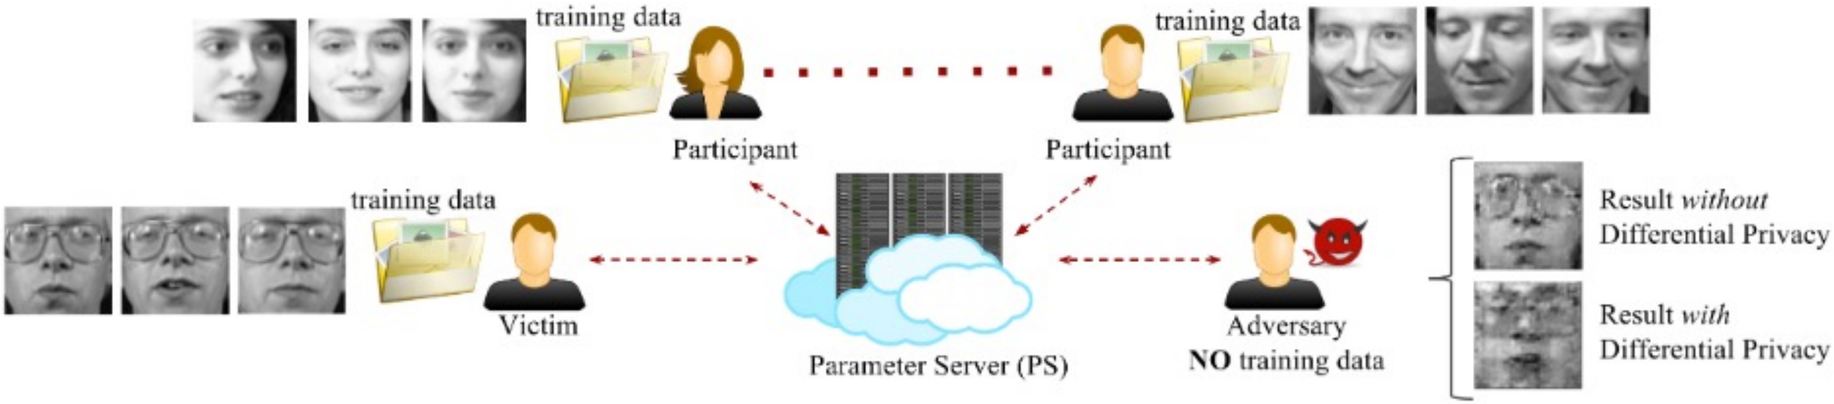

Figure 7: Collaborative deep learning with 41 participants. All 40 honest users train their respective models on distinct faces. The adversary has no local data. The GAN on the adversary's device is able to reconstruct the face stored on the victim's device (even when DP is enabled).

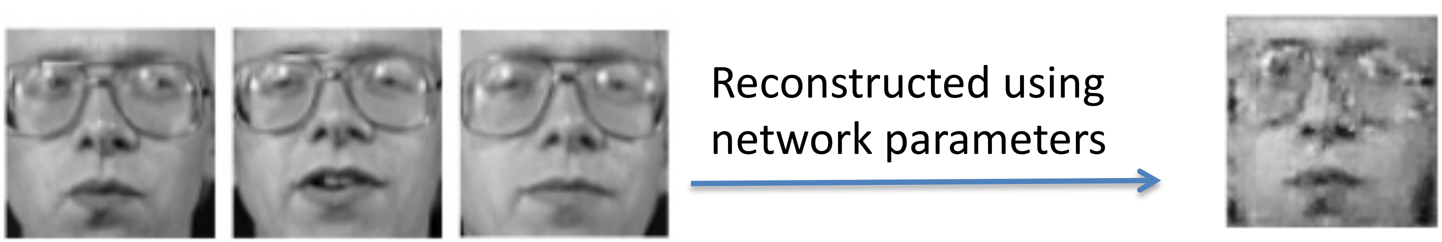

Deep Models Under the GAN: Information Leakage from Collaborative Deep Learning, Hitaj et al., ACM CCS'17

Federated Learning with Differential Privacy

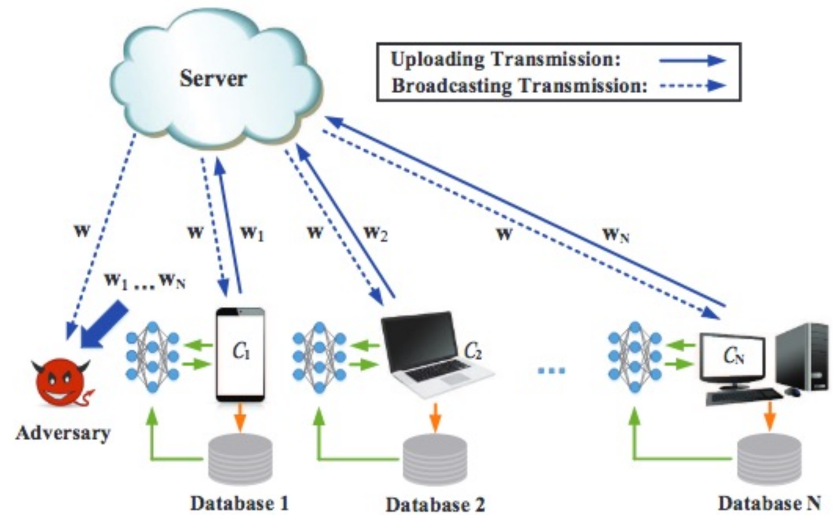

Figure 1: A FL training model with hidden adversaries who can eavesdrop trained parameters from both the clients and the server.

Algorithm 1: Noising before Aggregation FL

Data: T, $\mathbf{w}^{(0)}$, μ, ϵ and δ

Initialization: $t = 1$ and $\mathbf{w}_i^{(0)} = \mathbf{w}^{(0)}, \forall i$

while $t \leq T$ **do**

- **Local training process:**
- **while** $C_i \in \{C_1, C_2, \ldots, C_N\}$ **do**
  - Update the local parameters $\mathbf{w}_i^{(t)}$ as

    $$\mathbf{w}_i^{(t)} = \arg\min_{\mathbf{w}_i} \left(F_i(\mathbf{w}_i) + \frac{\mu}{2} \|\mathbf{w}_i - \mathbf{w}^{(t-1)}\|^2 \right)$$
  - Clip the local parameters

    $$\mathbf{w}_i^{(t)} = \mathbf{w}_i^{(t)} / \max\left(1, \frac{\|\mathbf{w}_i^{(t)}\|}{C}\right)$$
  - Add noise and upload parameters

    $$\widetilde{\mathbf{w}}_i^{(t)} = \mathbf{w}_i^{(t)} + \mathbf{n}_i^{(t)}$$
- **Model aggregating process:**
- Update the global parameters $\mathbf{w}^{(t)}$ as

  $$\mathbf{w}^{(t)} = \sum_{i=1}^{N} p_i \widetilde{\mathbf{w}}_i^{(t)}$$
- The server broadcasts global noised parameters

  $$\widetilde{\mathbf{w}}^{(t)} = \mathbf{w}^{(t)} + \mathbf{n}_{\mathrm{D}}^{(t)}$$
- **Local testing process:**
- **while** $C_i \in \{C_1, C_2, \ldots, C_N\}$ **do**
  - Test the aggregating parameters $\widetilde{\mathbf{w}}^{(t)}$ using local dataset
- $t \leftarrow t + 1$

Result: $\widetilde{\mathbf{w}}^{(T)}$

K. Wei et al., "Federated Learning With Differential Privacy: Algorithms and Performance Analysis," in IEEE Transactions on Information Forensics and Security, vol. 15, pp. 3454-3469, 2020.

Federated Learning with Differential Privacy

- Three key properties
  - There is a tradeoff between convergence performance and privacy protection levels, i.e., better convergence performance leads to a lower protection level
  - Given a fixed privacy protection level, increasing the number N of overall clients participating in FL can improve the convergence performance
  - There is an optimal number aggregation times (communication rounds) in terms of convergence performance for a given protection level

K. Wei et al., "Federated Learning With Differential Privacy: Algorithms and Performance Analysis," in IEEE Transactions on Information Forensics and Security, vol. 15, pp. 3454-3469, 2020.

Split Learning Network (SplitNN)

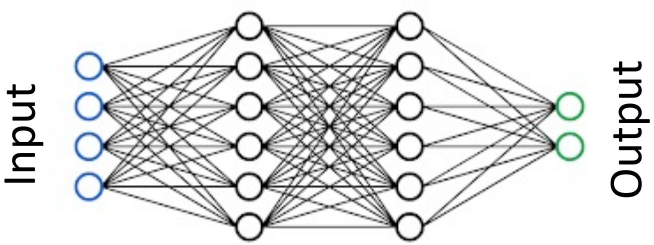

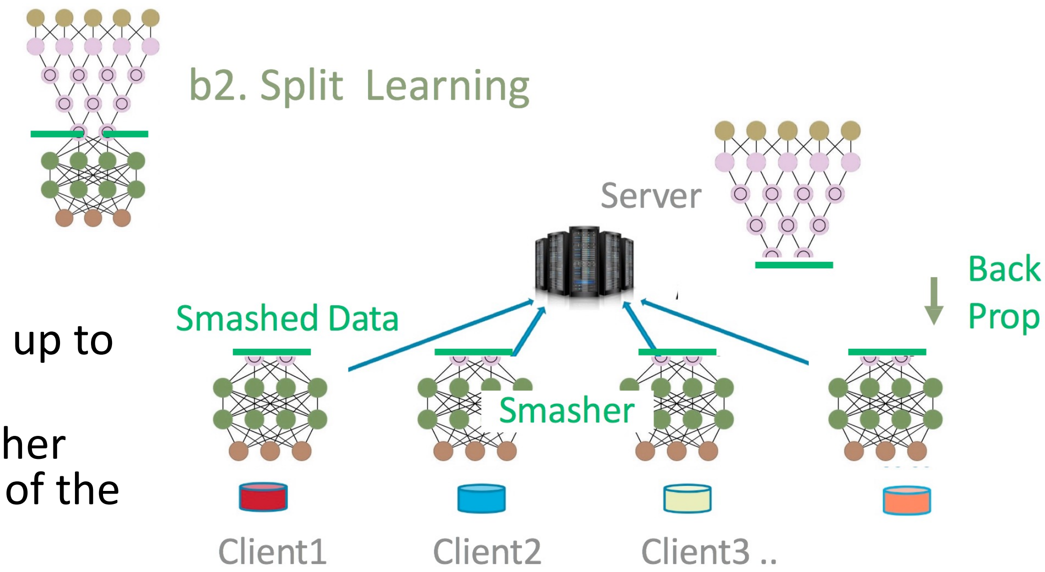

- Each client trains a partial deep network up to a specific layer (cut layer)
- Outputs at the cut layer are sent to another entity (server) which completes the rest of the training
- The gradients are now back propagated again from its last layer until the cut layer in a similar fashion
- The gradients at the cut layer are sent back to client centers
- This process is continued until the distributed split learning network is trained
- Computational, communication, and memory efficient
- Large number of clients: Split learning shows positive results

Gupta, Otkrist and Raskar, Ramesh, *Distributed learning of deep neural network over multiple agents*, Journal of Network and Computer Applications, Vol.116, pp.1–8, 2018.

https://splitlearning.github.io/

Image credit: Raskar MIT

Federated Learning - Tools

- OpenMind (www.openmined.org)
  - An open-source community whose goal is to make the world more privacy-preserving by lowering the barrier-to-entry to private AI technologies.
- PySyft: Python library for secure and private Deep Learning
  - https://github.com/OpenMined/PySyft)
- TensorFlow Federated
  - Machine learning on decentralized data
  - https://www.tensorflow.org/federated
- Federated-Learning (PyTorch)
  - https://github.com/AshwinRJ/Federated-Learning-PyTorch

Applications

- Face recognition
- Face presentation attack detection
  - Multi-institutional collaboration
- Mobile-based active authentication
  - Learning over smart phones
- Thermal to visible face synthesis

Federated Face Recognition

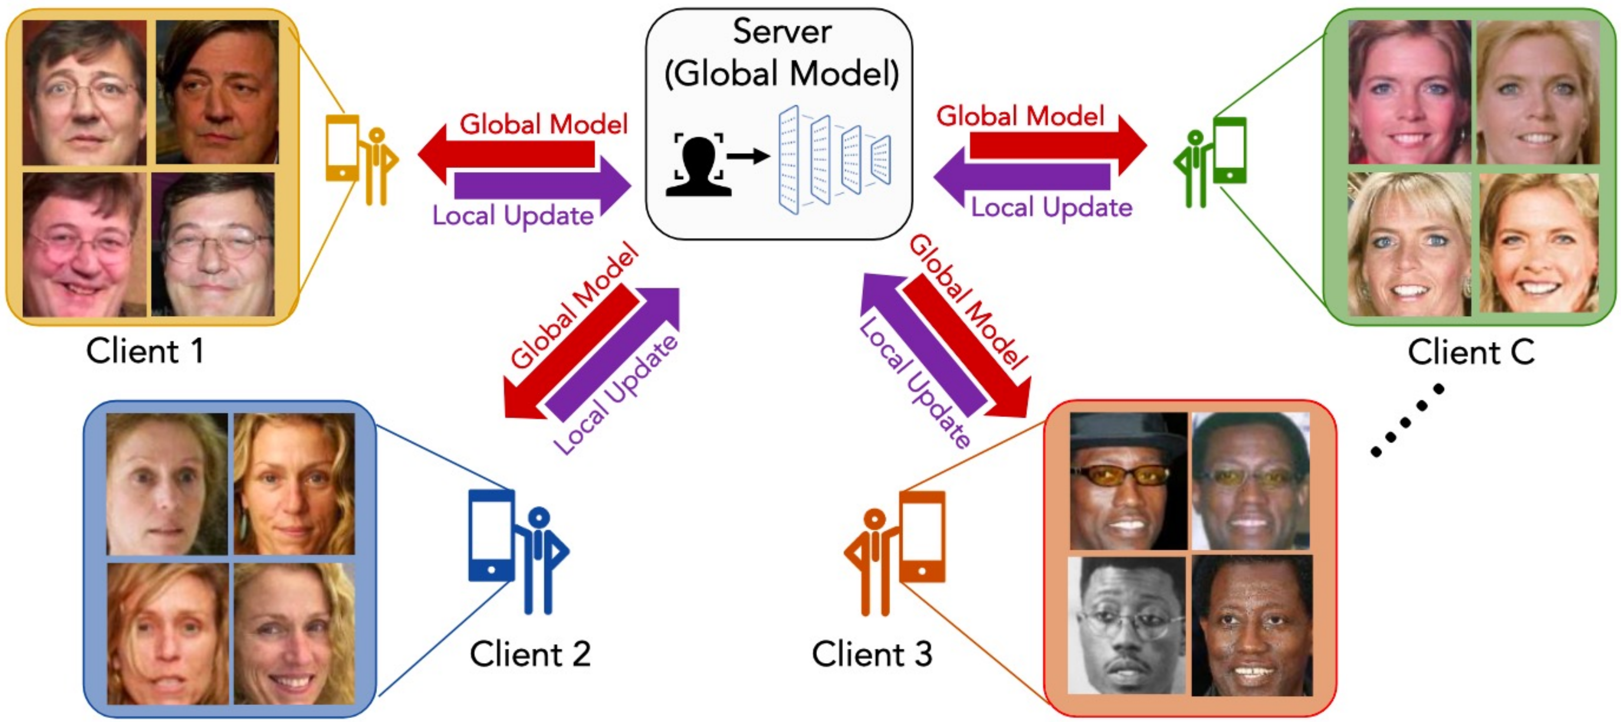

- Learning over smartphones
- One identity per client
  - Learning with only positive labels

[1] FedFace: Collaborative Learning of Face Recognition Model, D. Aggarwal et al. arXiv 2021.
https://arxiv.org/pdf/2104.03008.pdf
[2] Federated Learning with Only Positive Labels, F. X. Yu et al. ICML 2020.
https://arxiv.org/pdf/2004.10342.pdf

Federated Face Recognition

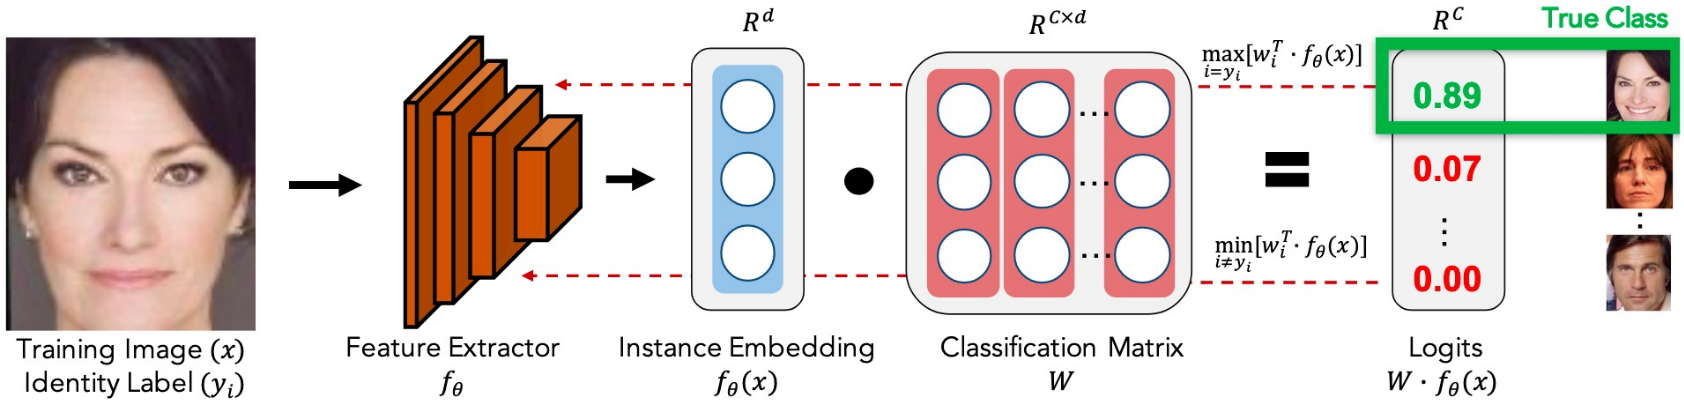

Figure 2: An overview of the training framework used by prevailing DNN-based AFR systems. An input x with a label y_i is passed through a feature extractor f_θ to obtain the feature vector $f_\theta(x)$. The feature vector is then multiplied with the classification matrix W to get the logits or the likelihood of x belonging to each of the C identities. We then maximize the similarity between the feature vector and the positive class embedding w_i and minimize the similarity between the feature vector and negative class embeddings (Red line indicates back-propagation of the loss through the model). In the FL setup, since each client does not have access to class embeddings of other clients/identities, the client cannot minimize the second term of the training objective.

$$l(x,y)^4 = \alpha \cdot (d(f_{\theta_t}(x), w_y))^2 + \beta \cdot \sum_{c\neq y} (\max\{0, v - d(f_{\theta_t}(x), w_c)\})^2$$

Similar to pairwise ranking loss

[1] FedFace: Collaborative Learning of Face Recognition Model, D. Aggarwal et al. arXiv 2021. https://arxiv.org/pdf/2104.03008.pdf
[2] Federated Learning with Only Positive Labels, F. X. Yu et al. ICML 2020. https://arxiv.org/pdf/2004.10342.pdf

Spreadout Regularization

- Impose a geometric regularizer (Spreadout regularizer [2]) after each round to encourage classes to be spreadout in the embedding space.

$$l_{pos}(f_{\theta_t}(x), i) = \max\left(0, m - (w_t^i)^T f_{\theta_t}(x)\right)^2$$

$$W_t = [w_t^1, w_t^2, \ldots, w_t^C]^T$$

$$reg_{sp}(W_t)^{5} = \sum_{c \in [C]} \sum_{\hat{c} \neq c} (\max\{0, v - d(w_t^c, w_t^{\hat{c}})\})^2$$

Encourages classes to be spreadout in the embedding space.

[1] FedFace: Collaborative Learning of Face Recognition Model, D. Aggarwal et al. arXiv 2021.
https://arxiv.org/pdf/2104.03008.pdf
[2] Federated Learning with Only Positive Labels, F. X. Yu et al. ICML 2020.
https://arxiv.org/pdf/2004.10342.pdf

FedFace

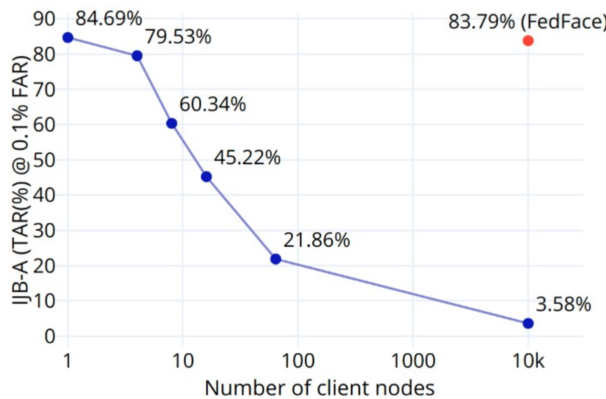

Figure 4: Effect of the number of clients on the FedAvg [25] algorithm. We divide the 10,000 subjects in CASIA-WebFace [33] equally into different client nodes for training. One client node denotes the conventional (non-federated) way of training AFR systems while the $10k$ clients represents the problem we are tackling with face images of one identity per client. We evaluate on IJB-A [20]. Note that the x-axis is in log scale.

Trained on CASIA-WebFace

Table 1: Face Verification performance of *FedFace* on standard face recognition benchmarks LFW [17], IJB-A [20] and IJB-C [24]. We use CosFace [30] (64-layer) as our feature extractor.

Method	**Training Data (Centrally aggregated or distributed)**	**LFW [17]**	**IJB-A [20]**	**IJB-C [24]**
		LFW Accuracy(%)	**TAR @ 0.1% FAR**	**TAR @ 0.1% FAR**
Baseline	Centrally aggregated	99.15%	81.43%	84.78%
Fine-tuning baseline in a non-federated manner	Centrally aggregated	99.32%	84.18%	88.76%
Randomly Initialized class embeddings	Distributed	94.61%	70.13%	69.30%
Proposed FedFace	**Distributed**	**99.28%**	**83.79%**	**88.21%**

FedFace: Collaborative Learning of Face Recognition Model, D. Aggarwal et al. arXiv 2021.
https://arxiv.org/pdf/2104.03008.pdf

Federated Face Presentation Attack Detection (FedPAD)

Marcel et al.

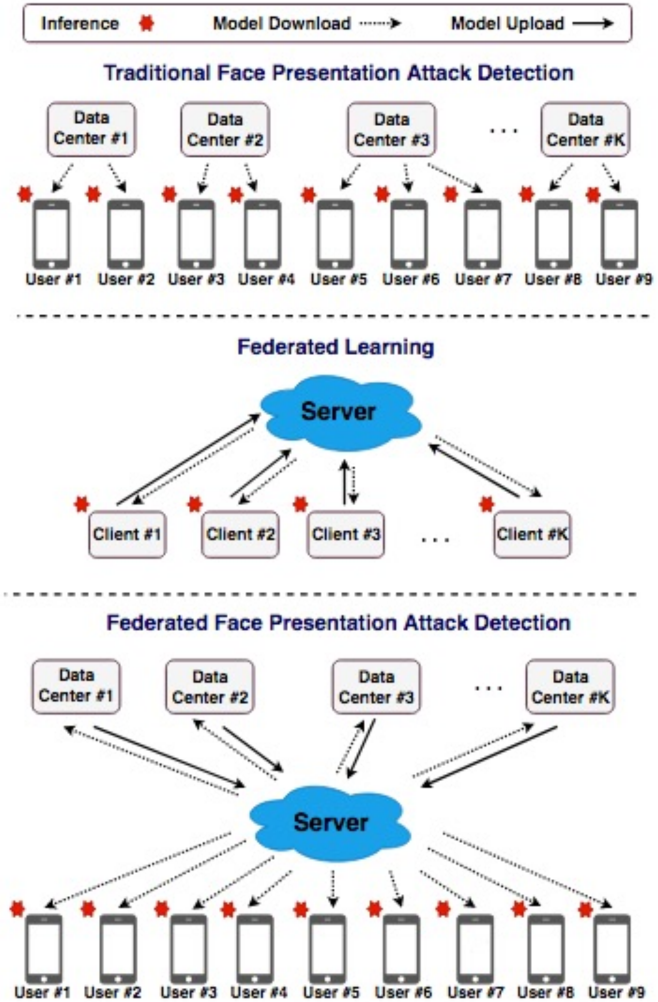

Figure 1. Comparison between fPAD (top), traditional federated learning (middle) and the proposed FedPAD (bottom). FedPAD can be a regarded as a special case of traditional federated learning.

Shao et al, FG 2021
https://arxiv.org/pdf/2005.14638.pdf

FedPAD Framework

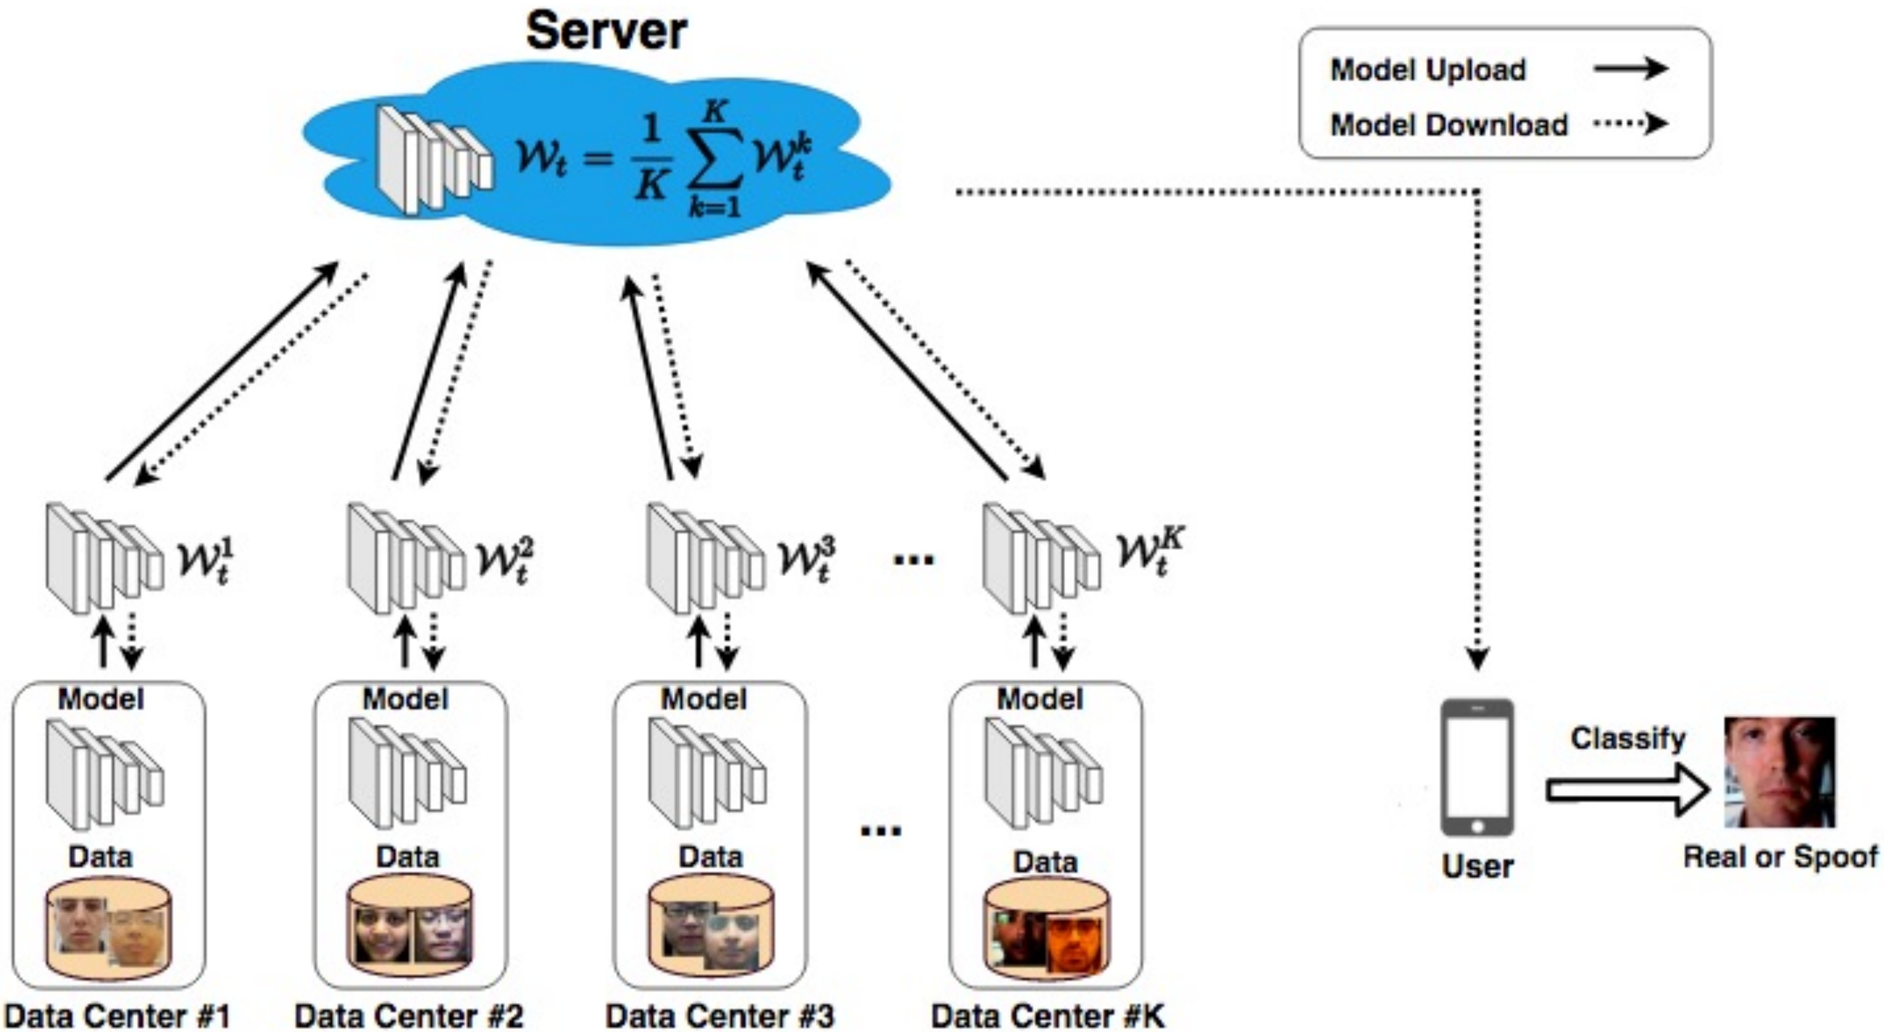

Shao et al, 2020
https://arxiv.org/pdf/2005.14638.pdf

FedPAD Data

Table 1. Comparison of seven experimental datasets.

Dataset	Extra light	Complex background	Attack type	Display devices
C	No	Yes	Printed photo Cut photo Replayed video	iPad
I	Yes	Yes	Printed photo Display photo Replayed video	iPhone 3GS iPad
M	No	Yes	Printed photo Replayed video	iPad Air iPhone 5S
O	Yes	No	Printed photo Display photo Replayed video	Dell 1905FP Macbook Retina
S	Yes	Yes	Printed photo Display photo Replayed video	Dell 1905FP iPad Pro iPhone 7 Galaxy S8 Asus MB168B
3	No	No	Thatsmyface 3D mask	Kinect
H	Yes	Yes	Thatsmyface 3D mask REAL-f mask	MV-U3B

FedPAD Results

Table 2. Comparison with models trained by data from single data center and various data centers.

Methods	Data Centers	User	HTER (%)	EER (%)	AUC (%)	Avg. HTER	Avg. EER	Avg. AUC
Single	O	M	41.29	37.42	67.93	36.43	34.31	70.36
	C	M	27.09	24.69	82.91			
	I	M	49.05	20.04	85.89			
	O	C	31.33	34.73	73.19			
	M	C	39.80	40.67	66.58			
	I	C	49.25	47.11	55.41			
	O	I	42.21	43.05	54.16			
	C	I	45.99	48.55	51.24			
	M	I	48.50	33.70	66.29			
	M	O	29.80	24.12	84.86			
	C	O	33.97	21.24	84.33			
	I	O	46.95	35.16	71.58			
Fused	O&C&I	M	34.42	23.26	81.67	35.75	31.29	73.89
	O&M&I	C	38.32	38.31	67.93			
	O&C&M	I	42.21	41.36	59.72			
	I&C&M	O	28.04	22.24	86.24			
Ours	O&C&I	M	19.45	17.43	90.24	**32.17**	**28.84**	**76.51**
	O&M&I	C	42.27	36.95	70.49			
	O&C&M	I	32.53	26.54	73.58			
	I&C&M	O	34.44	34.45	71.74			
All (Upper Bound)	O&C&I	M	21.80	17.18	90.96	27.26	25.09	80.42
	O&M&I	C	29.46	31.54	76.29			
	O&C&M	I	30.57	25.71	72.21			
	I&C&M	O	27.22	25.91	82.21			

Single: Obtain a trained model from one data center.
Fused: Obtain multiple trained models from several data centers and fuse their prediction scores during inference
Ours: Performance of a trained model is evaluated against a dataset that has not been observed during training
All: Model is trained with data from all available data centers (not privacy preserving)

FedPAD Results

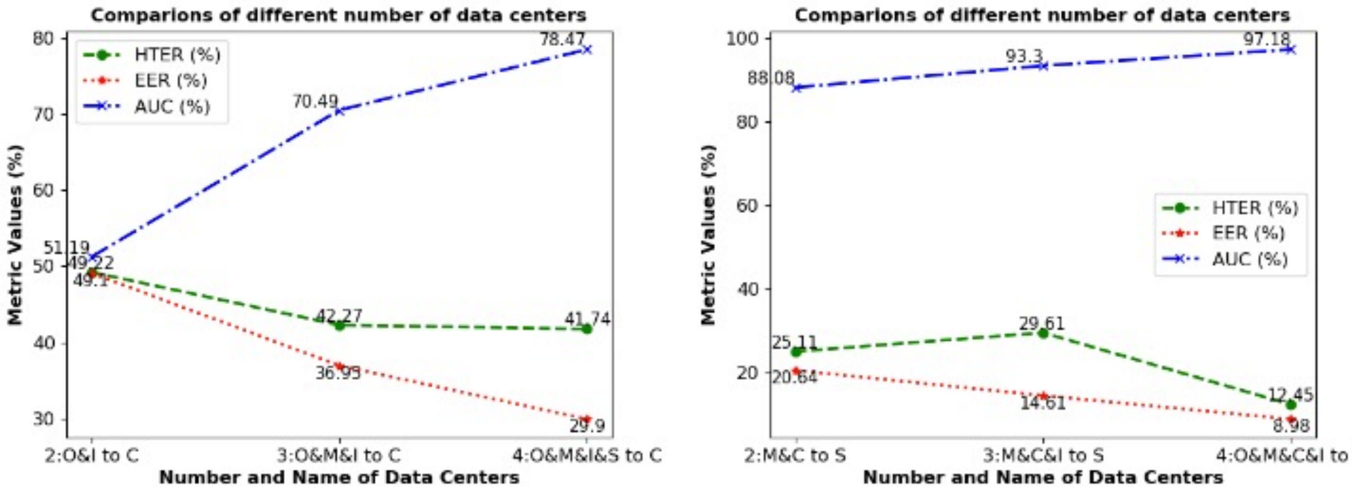

Figure 5. Comparison of different number of data centers.

Table 3. Effect of using different types of spoof attacks

Methods	Data Centers	User	HTER (%)	EER (%)	AUC (%)
Single	I (Print)	M (Print, Video)	38.82	33.63	72.46
	O (Video)	M (Print, Video)	35.76	28.55	78.86
Fused	I (Print) & O (video)	M (Print, Video)	35.22	**25.56**	81.54
Ours	I (Print) & O (video)	M (Print, Video)	**30.51**	26.10	**84.82**

Table 4. Impact of adding data centers with diverse attacks

Data Centers	User	HTER (%)	EER (%)	AUC (%)
O&C&I&M (2D)	H (3D)	47.02	18.31	85.06
O&C&I&M (2D)&3 (3D)	H (3D)	**34.70**	**14.20**	**92.35**

Test-Time Adaptive FedPAD

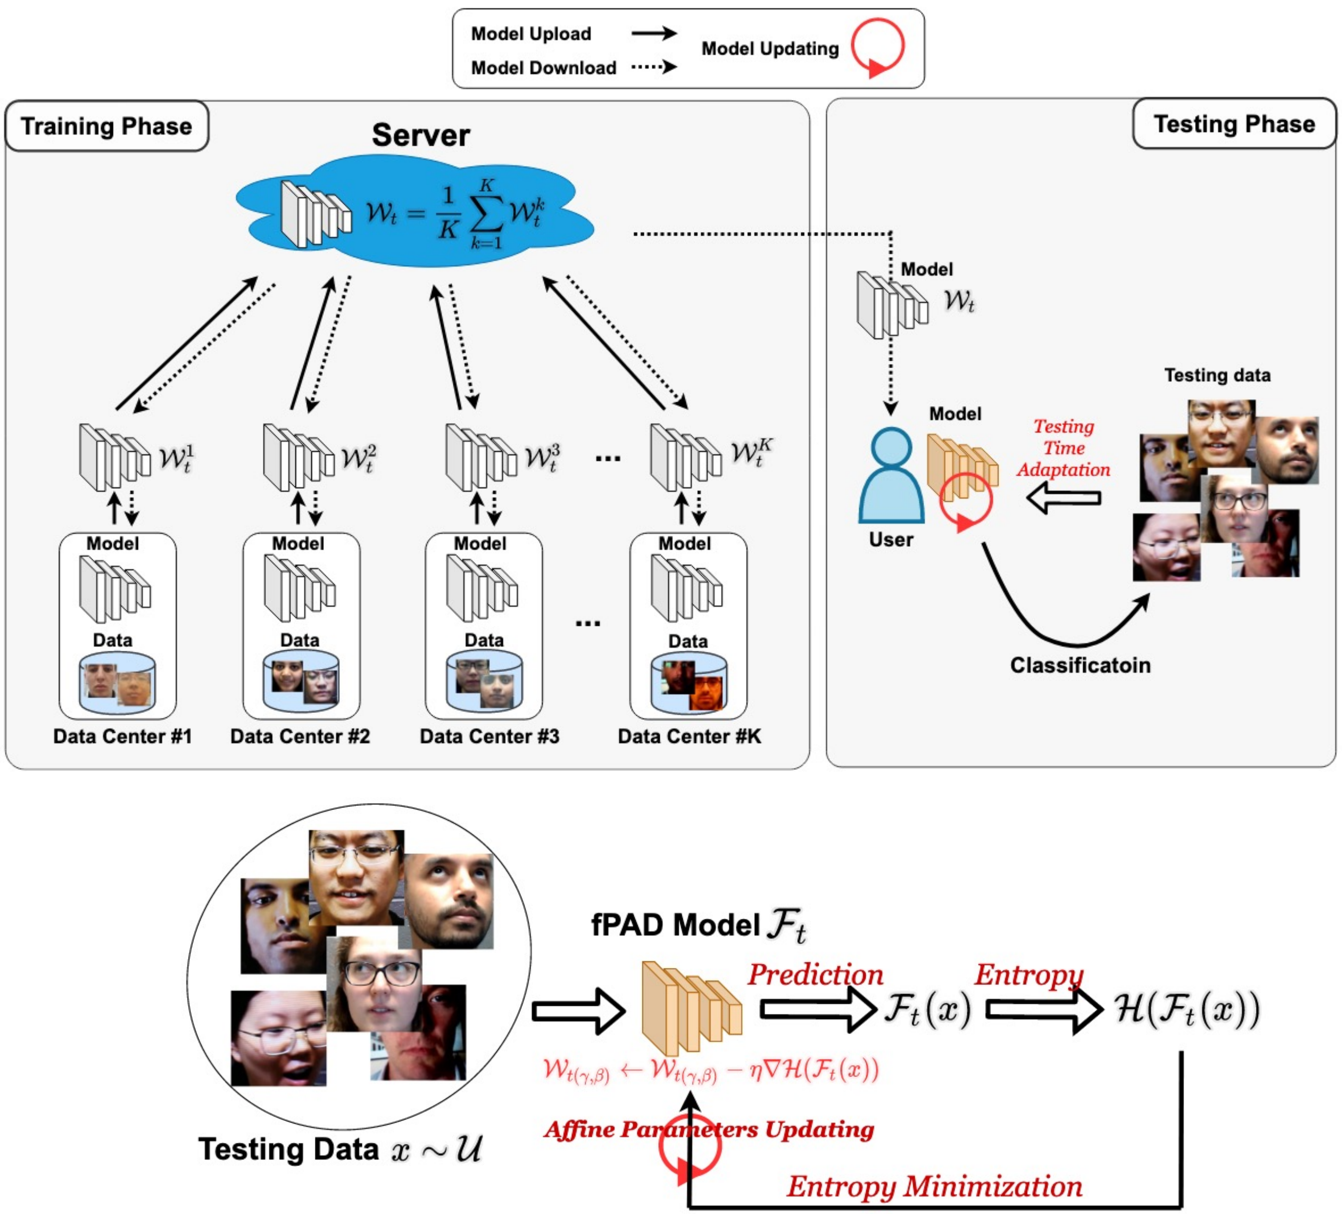

Test-Time Adaptive FedPAD Results

COMPARISON WITH MODELS TRAINED BY DATA FROM SINGLE DATA CENTER AND VARIOUS DATA CENTERS.

Methods	Data Centers	User	HTER (%)	EER (%)	AUC (%)	Avg. HTER	Avg. EER	Avg. AUC
Single	O	M	41.29	37.42	67.93	41.61	36.66	67.07
	C	M	27.09	24.69	82.91			
	I	M	49.05	20.04	85.89			
	O	C	31.33	34.73	73.19			
	M	C	39.80	40.67	66.58			
	I	C	49.25	47.11	55.41			
	O	I	42.21	43.05	54.16			
	C	I	45.99	48.55	51.24			
	M	I	48.50	33.70	66.29			
	M	O	29.80	24.12	84.86			
	C	O	33.97	21.24	84.33			
	I	O	46.95	35.16	71.58			
Fused	O&C&I	M	34.42	23.26	81.67	35.75	31.29	73.89
	O&M&I	C	38.32	38.31	67.93			
	O&C&M	I	42.21	41.36	59.72			
	I&C&M	O	28.04	22.24	86.24			
FedPAD	O&C&I	M	19.45	17.43	90.24	32.17	28.84	76.51
	O&M&I	C	42.27	36.95	70.49			
	O&C&M	I	32.53	26.54	73.58			
	I&C&M	O	34.44	34.45	71.74			
All	O&C&I	M	21.80	17.18	90.96	27.26	25.09	80.42
	O&M&I	C	29.46	31.54	76.29			
	O&C&M	I	30.57	25.71	72.21			
	I&C&M	O	27.22	25.91	82.21			
Ours	O&C&I	M	14.70	16.64	90.57	**23.18**	**23.88**	**83.40**
	O&M&I	C	26.33	29.75	77.77			
	O&C&M	I	28.61	26.04	82.07			
	I&C&M	O	23.09	23.09	83.21			

Active Authentication (AA)

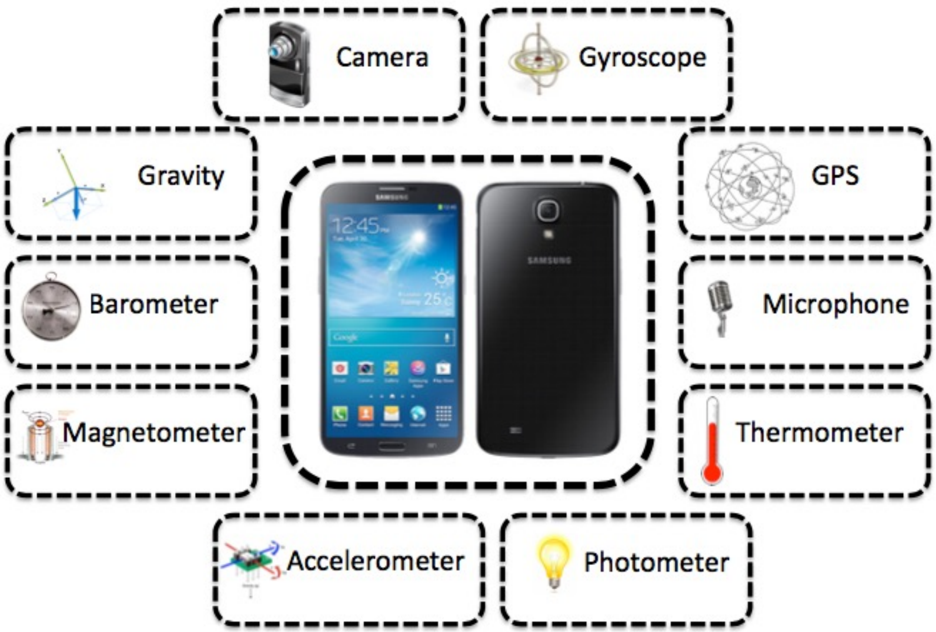

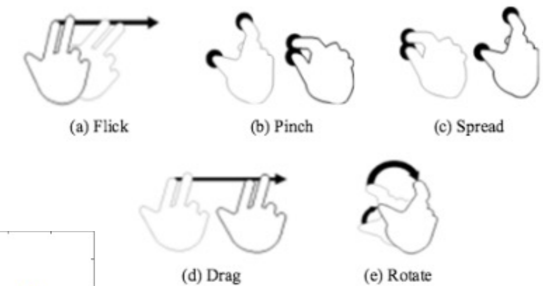

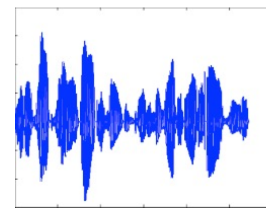

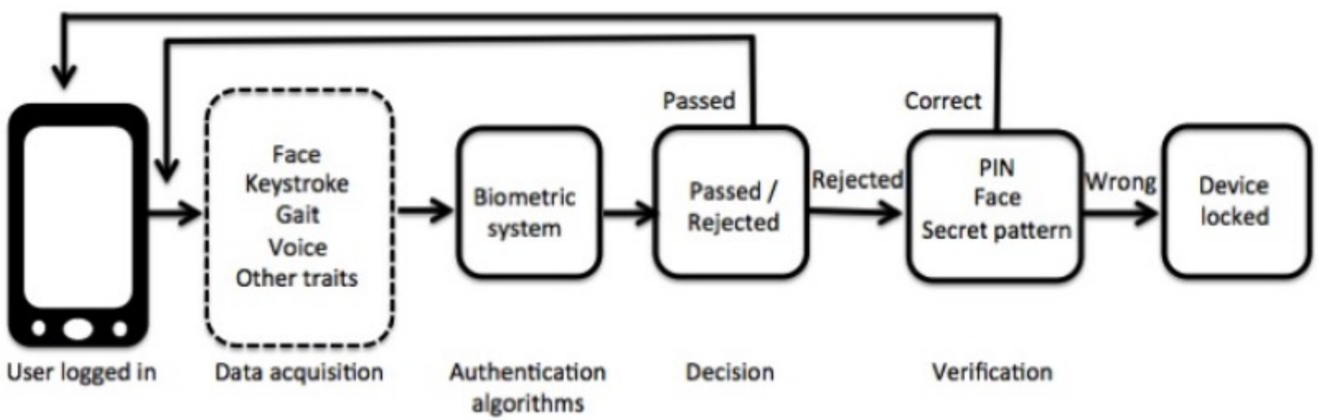

V. M. Patel, R. Chellappa, D. Chandra and B. Barbello, "Continuous User Authentication on Mobile Devices: Recent progress and remaining challenges," in IEEE Signal Processing Magazine, vol. 33, no. 4, pp. 49-61, July 2016.

AA - OCC Problem

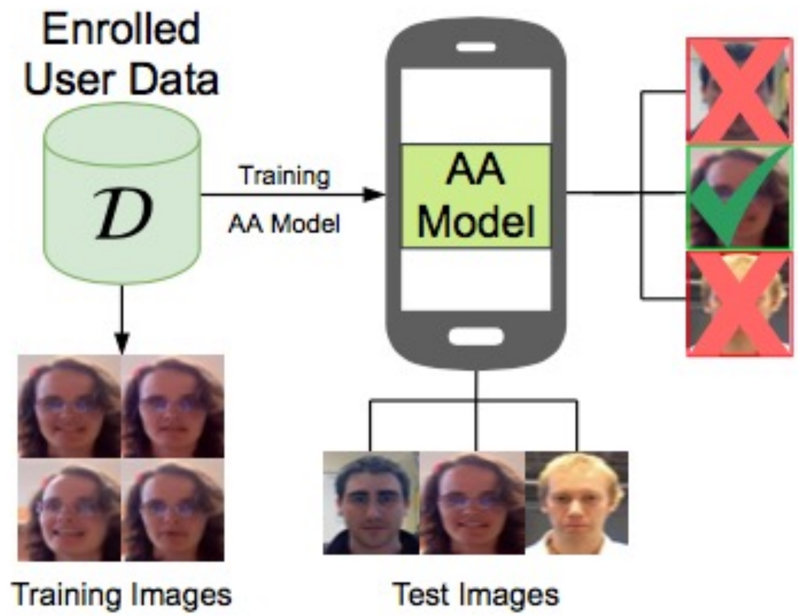

Federated AA Framework

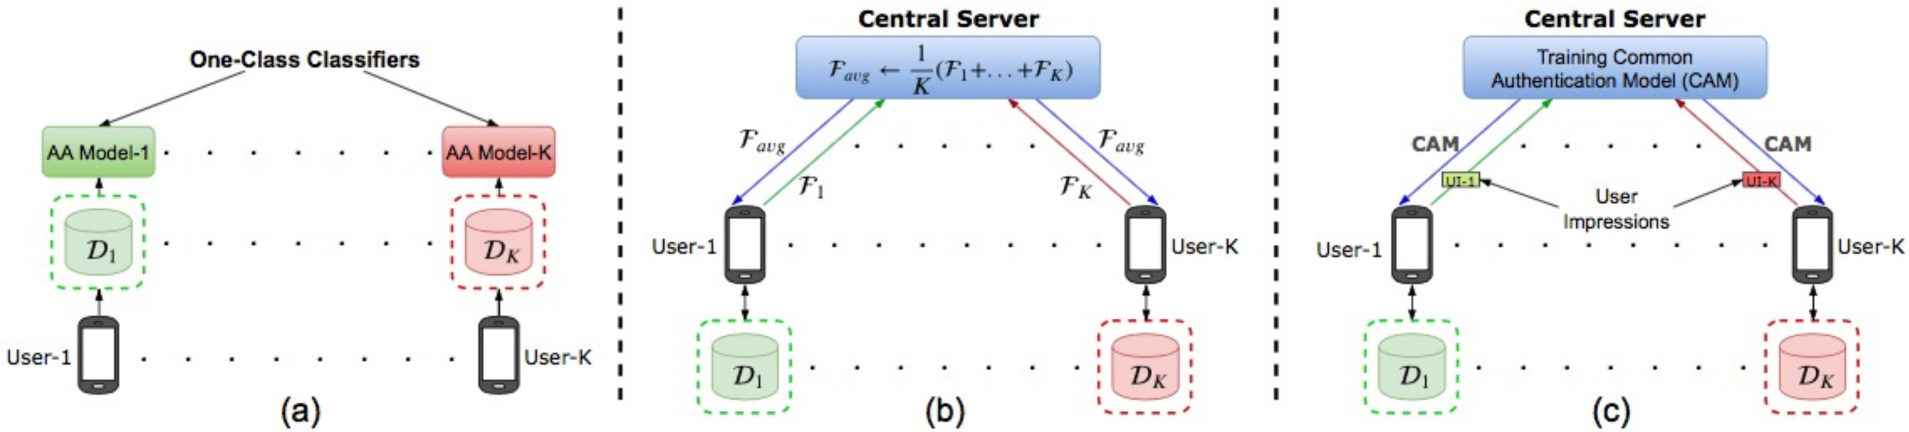

Figure 2. Active authentication based on (a) One class classification, (b) Federated Averaging, and (c) Proposed Method.

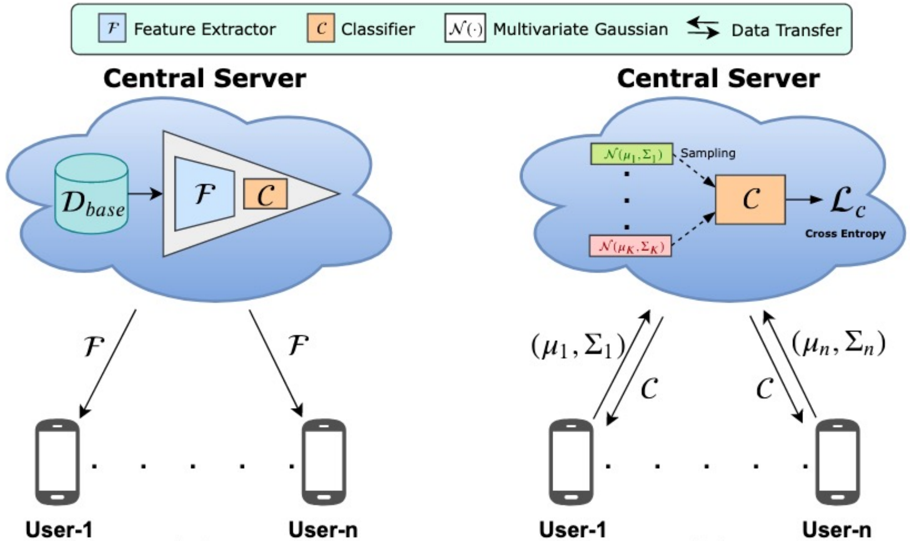

Federated AA

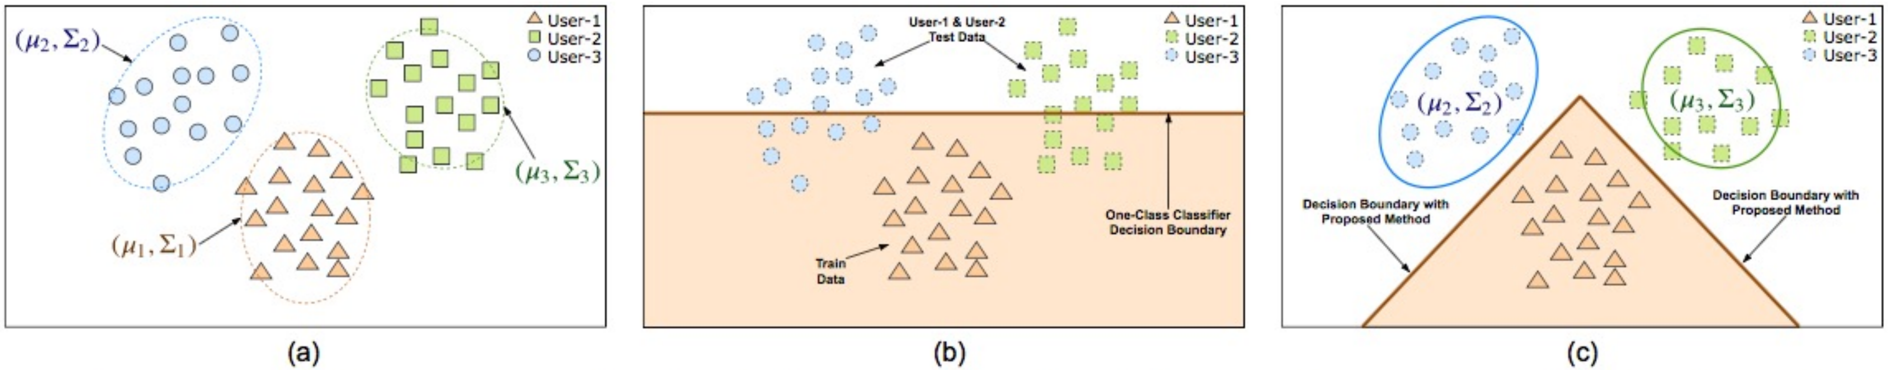

Figure 5. Toy example with three users to show the effectiveness of proposed method compared to one-class modeling based methods. (a) Feature space location (mean μ_i) and shape (variance Σ_i) estimated for each user. (b) Modeling as a one-class classification problem to learn a decision boundary for user-1. When such a model is tested there are many samples from user-2 and user-3 that are mis-classified as user-1. (c) Learning decision boundary using proposed method to train the authentication model for user-1 using user-1, user-2 and user-3's mean and variance. This model does not make the same mistake of mis-classifying user-2 and user-3 data as user-1 similar to one-class based method. As visible from the figure, the learned decision boundary is also better in comparison to one-class method.

Federated AA - Results

(a) MOBIO

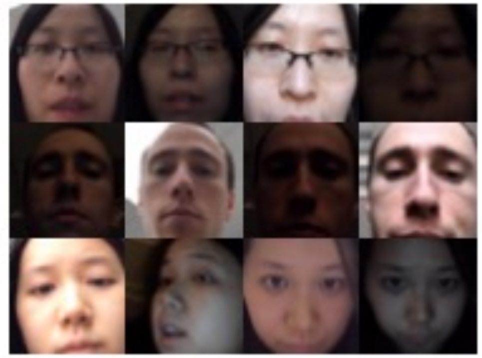

(b) UMDAA-01

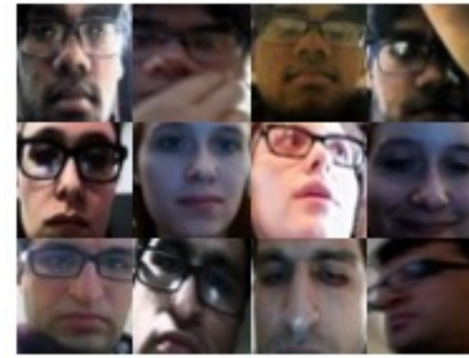

(c) UMDAA-02

Table 1. Performance comparison with state-of-the-art active authentication methods evaluated in terms of average detection accuracy. The best performing method for each dataset is shown in bold fonts.

	1SVM	k1SVM	SVDD	kSVDD	kNFST	1vSet	1MPM	DMPM	OC-ACNN	Proposed
MOBIO	0.632 (0.004)	0.748 (0.004)	0.582 (0.007)	0.763 (0.013)	0.560 (0.003)	0.670 (0.005)	0.768 (0.003)	0.825 (0.007)	0.938 (0.005)	**0.998 (0.003)**
UMDAA-01	0.622 (0.002)	0.731 (0.009)	0.615 (0.018)	0.701 (0.009)	0.567 (0.012)	0.593 (0.017)	0.816 (0.003)	0.869 (0.001)	0.891 (0.002)	**0.954 (0.005)**
UMDAA-02	0.614 (0.008)	0.649 (0.004)	0.515 (0.007)	0.550 (0.007)	0.556 (0.003)	0.538 (0.003)	0.722 (0.006)	0.760 (0.007)	0.735 (0.009)	**0.813 (0.006)**

Federated AA - Results

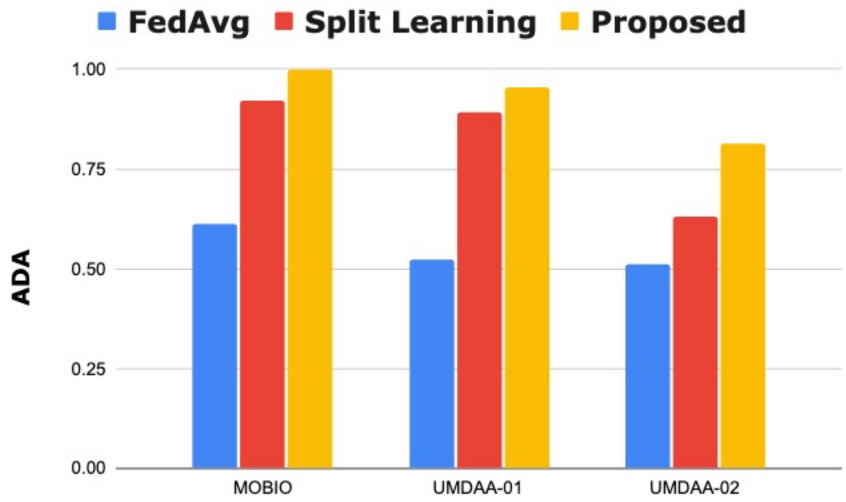

Federated Thermal to Visible Synthesis

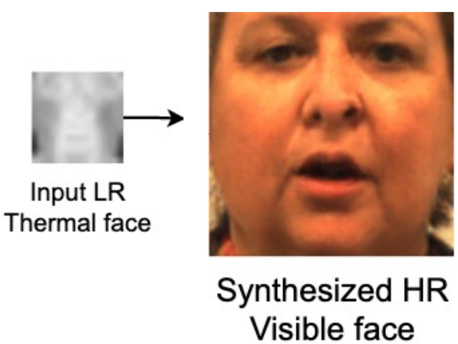

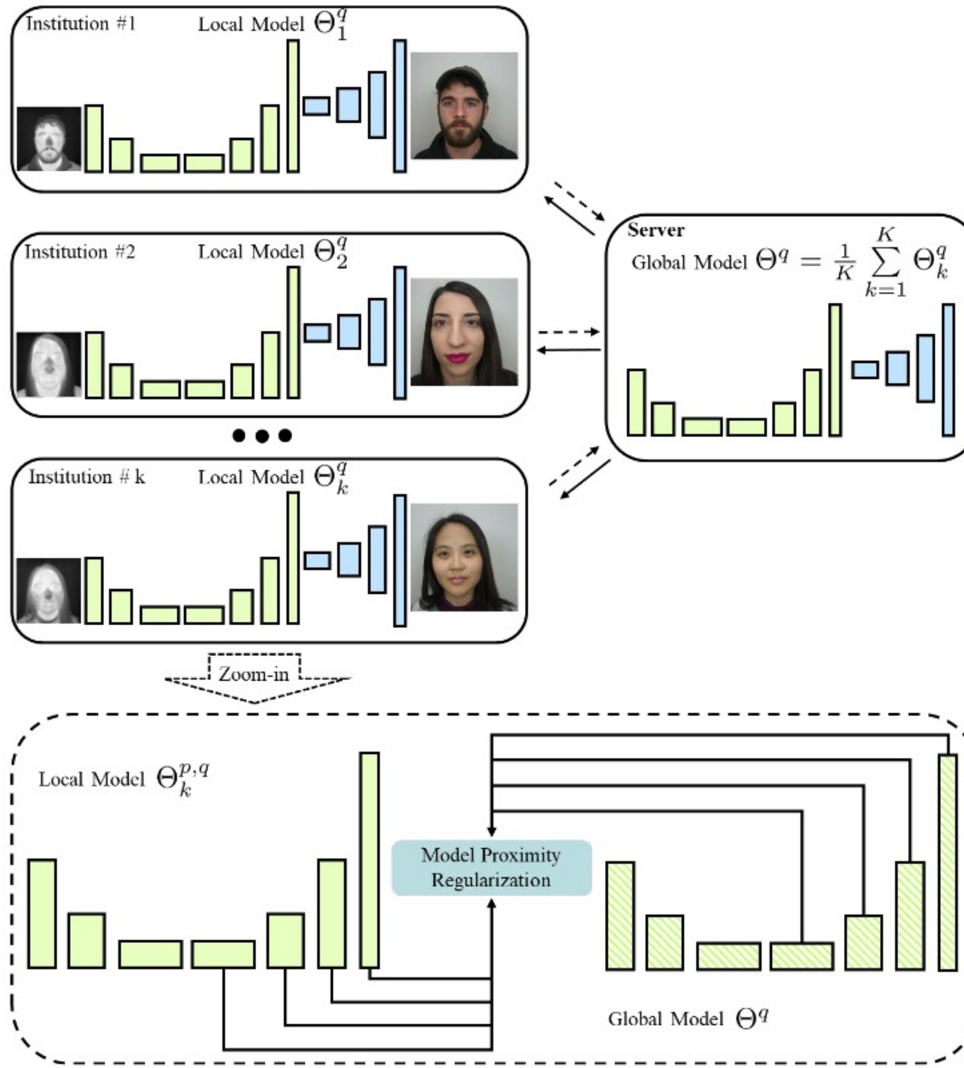

Mei, Guo & Patel, CVPR 2022

Federated Thermal to Visible Synthesis Results

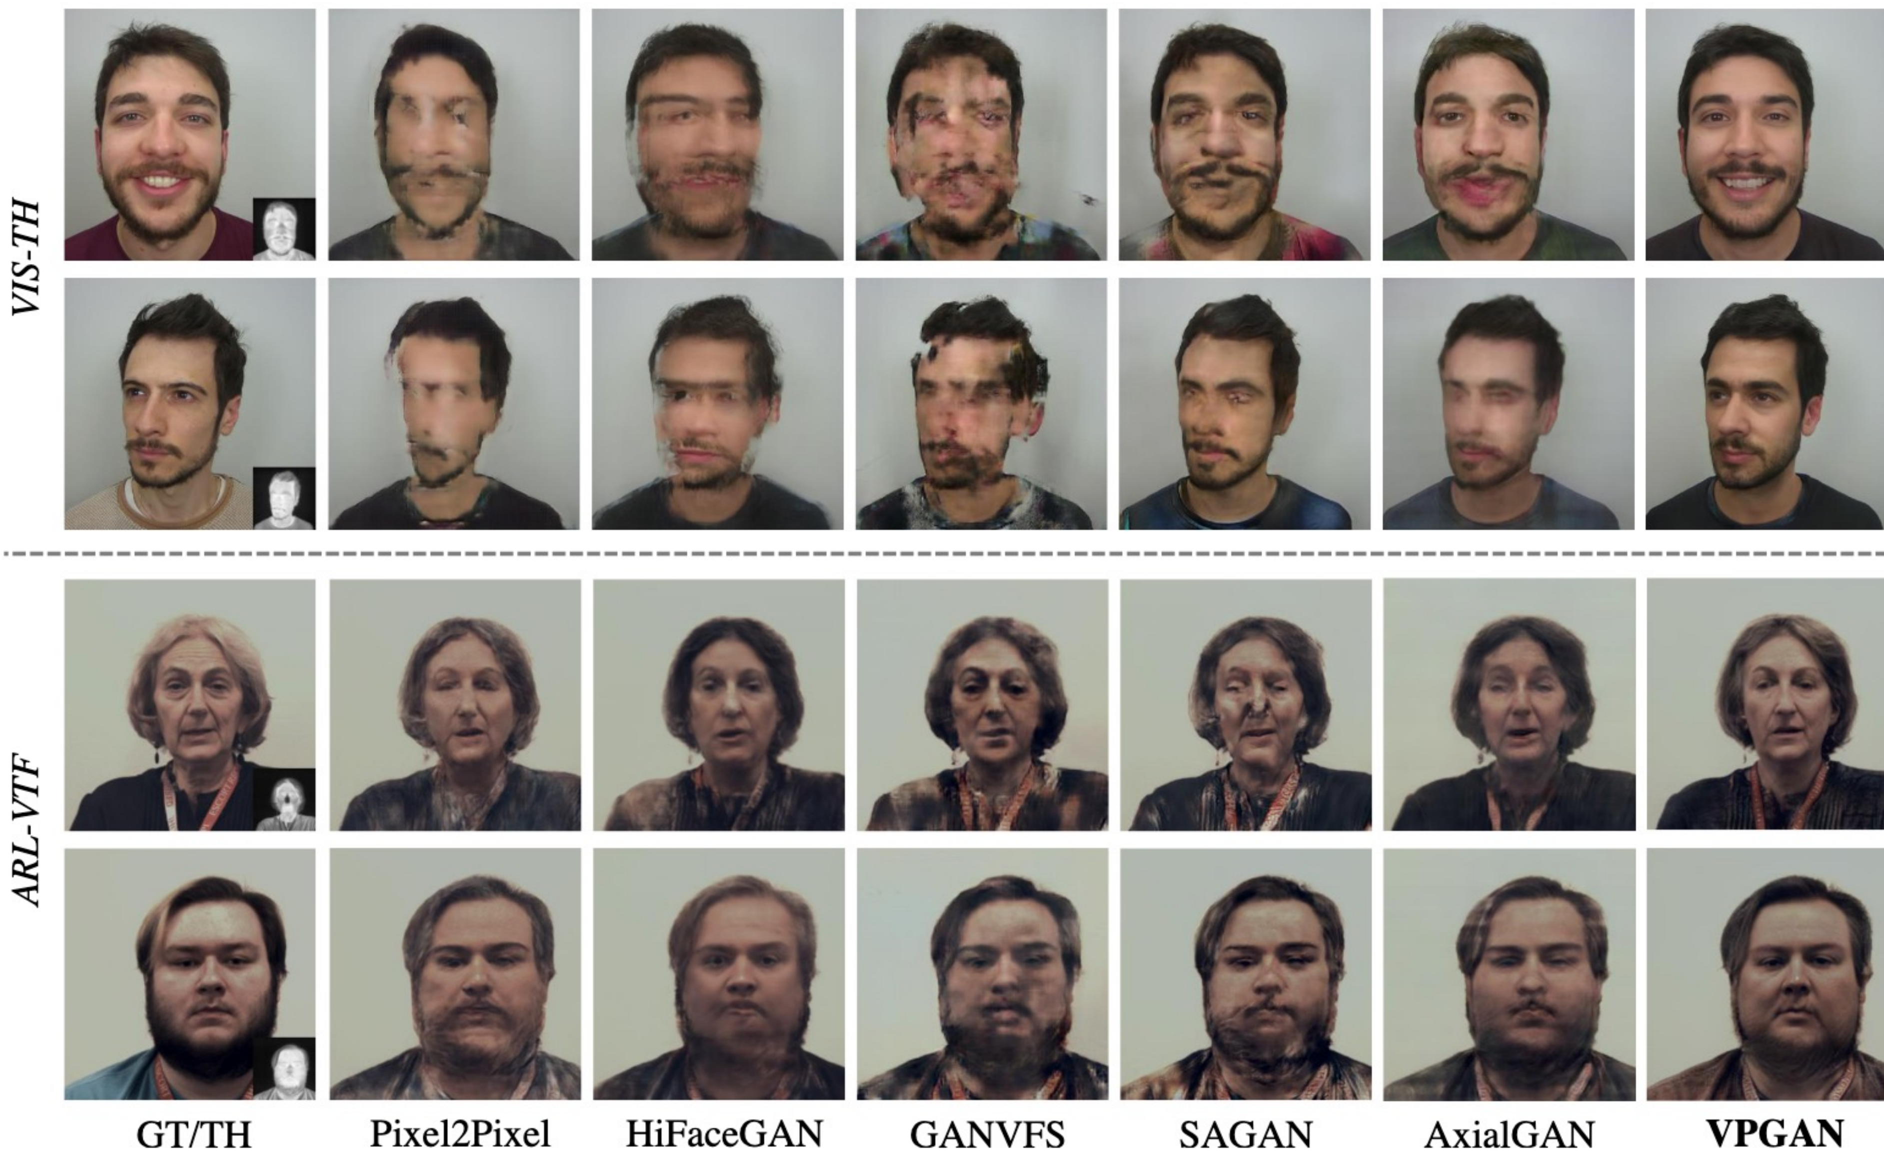

128 → 512

Federated Thermal to Visible Synthesis Results

Table 2. Verification results on the **VIS-TH** dataset.

Method	Rank-1	VR@FAR=1%	VR@FAR=0.1%
LightCNN [60]	30.48	8.57	2.86
Pixel2Pixel [19]	15.24	2.21	0.07
HiFaceGAN [62]	44.76	10.95	2.86
GANVFS [67]	18.11	7.29	1.90
SAGAN [6]	63.33	23.81	17.62
AxialGAN [18]	66.67	24.76	13.81
VPGAN (ours)	**76.67**	**45.71**	**20.00**

Table 4. Verification results on the **ARL-VTF** dataset.

Method	Rank-1	VR@FAR=1%	VR@FAR=0.1%
LightCNN [60]	11.07	9.24	4.57
Pixel2Pixel [19]	70.96	56.35	33.60
HiFaceGAN [62]	70.15	56.65	32.18
GANVFS [67]	70.76	45.99	22.03
SAGAN [6]	71.16	54.11	38.07
AxialGAN [18]	71.57	57.16	37.36
VPGAN (ours)	**74.16**	**59.96**	**41.27**

Table 3. Image quality results on the **ARL-VTF** dataset.

Methods	LPIPS↓	NIQE↓	Deg.↑	PSNR↑	SSIM↑
TH	0.6721	10.176	42.34	5.63	0.2940
Pixel2Pixel [19]	0.2038	6.298	70.67	19.46	0.7759
HiFaceGAN [62]	0.2166	7.274	70.11	19.67	**0.7954**
GANVFS [67]	0.2433	6.679	67.26	19.76	0.7511
SAGAN [6]	0.1925	6.155	71.12	20.11	0.7772
AxialGAN [18]	0.1998	6.223	69.75	20.17	0.7770
VPGAN (ours)	**0.1713**	**6.059**	**72.00**	**20.29**	0.7883

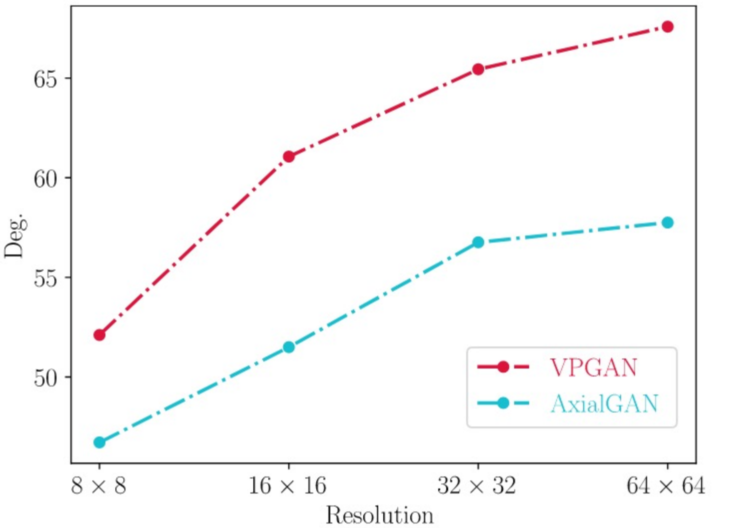

Figure 1. Effect of input resolution

Summary

- Federated learning promises to be an active area of research
- Open problems
  - Domain adaptive FL methods
  - Benchmarks
  - Unsupervised and semi-supervised FL
  - Privacy preserving FL methods
  - Novel FL models for biometrics and surveillance applications

Acknowledgments




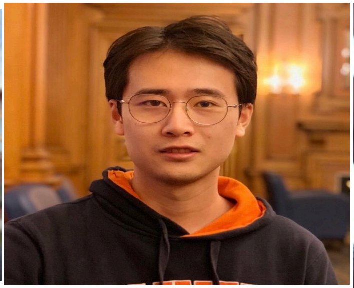

More Information,

Vision and Image Understanding (VIU) Lab @JHU

https://engineering.jhu.edu/vpatel36/

Thank You!