TMLR
Transactions on Machine Learning Research
Published June 2026

STEVENS INSTITUTE OF TECHNOLOGY 1870

FedIndex: Federated Domain Adaptation with Continuous Domain Indices

Qingyang Yu[1], Hao Wang[2], Qizhen Zhang[3], Hao Wang[1]

[1] Stevens Institute of Technology [2] Rutgers University [3] University of Toronto

IntelliSys Lab

Real-World Domain Shift Is Continuous

54.30%

Best categorical FDA on RotatingMNIST

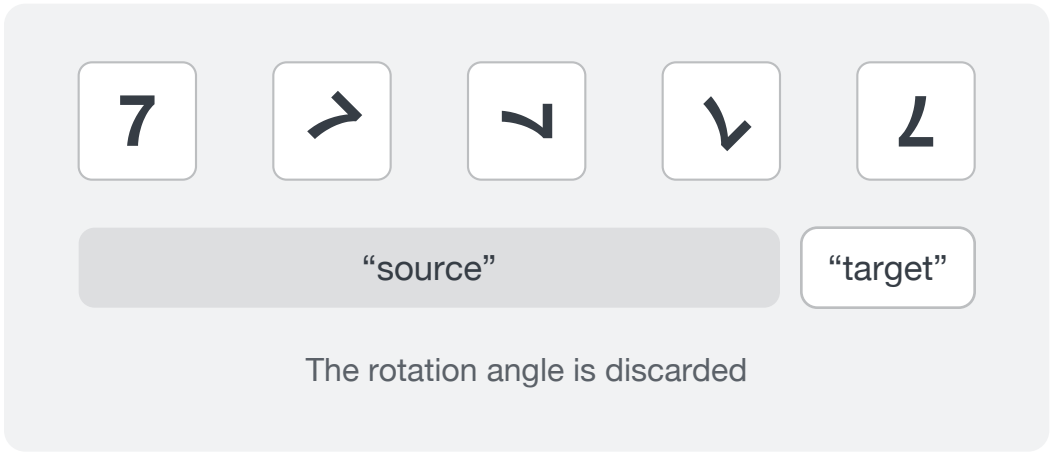

Categorical domain index

89.30%

Centralized CIDA on the same data

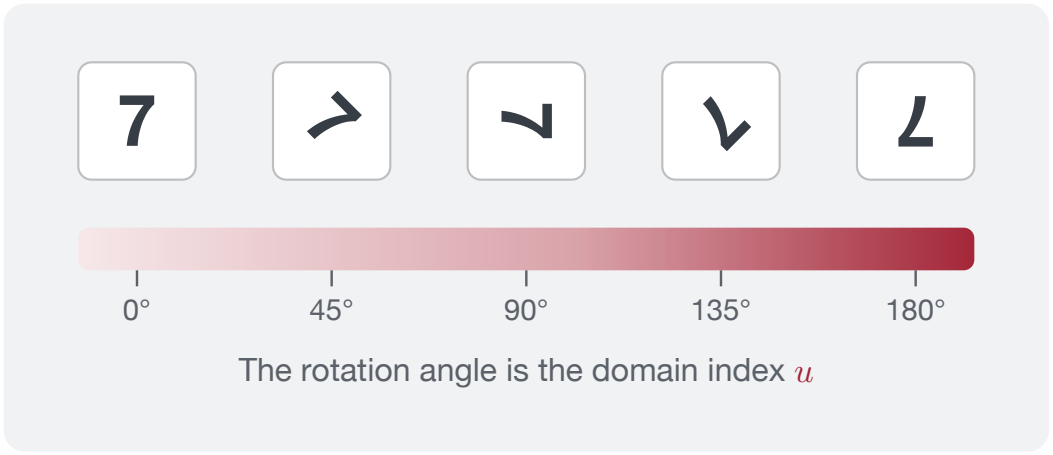

Continuous domain index

FDA Needs Continuous Domain Indices

[1] Q. Yu et al., “FedIndex: Federated Domain Adaptation with Continuous Domain Indices,” *TMLR*, 2026.
[2] H. Wang et al., “Continuously Indexed Domain Adaptation,” *ICML*, 2020.

Naively Federating CIDA Falls Short

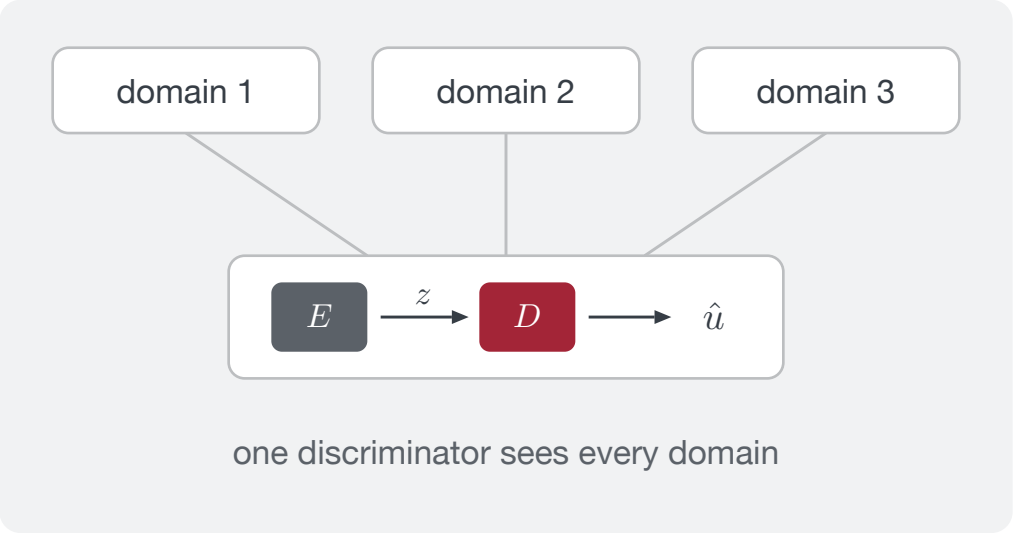

Centralized CIDA

One discriminator regresses the index over all domains

FedAvg + CIDA

Each local discriminator sees only its own client's domains

Averaging does not restore alignment.

[1] Q. Yu et al., "FedIndex: Federated Domain Adaptation with Continuous Domain Indices," *TMLR*, 2026.
[2] H. Wang et al., "Continuously Indexed Domain Adaptation," *ICML*, 2020.

Setting and Min–Max Objective

1 Index-invariant embeddings

Same distribution at every index

$$p(z \mid u_1) = p(z \mid u_2), \quad \forall\, u_1, u_2$$

2 Min–max game

Clients vs. one global discriminator

$$\min_{E_i, F_i} \max_{D} \quad V_p - \lambda_d V_d$$

3 Index regression loss

Squared error on the noised index

$$L_d = \big(D(z + \tau_i) - (u + \tau_u)\big)^2$$

First term: prediction loss on labeled source clients. Second term: index regression loss over all clients. τ: zero-mean Laplace noise.

[1] Q. Yu et al., "FedIndex: Federated Domain Adaptation with Continuous Domain Indices," *TMLR*, 2026.

One Round of FedIndex

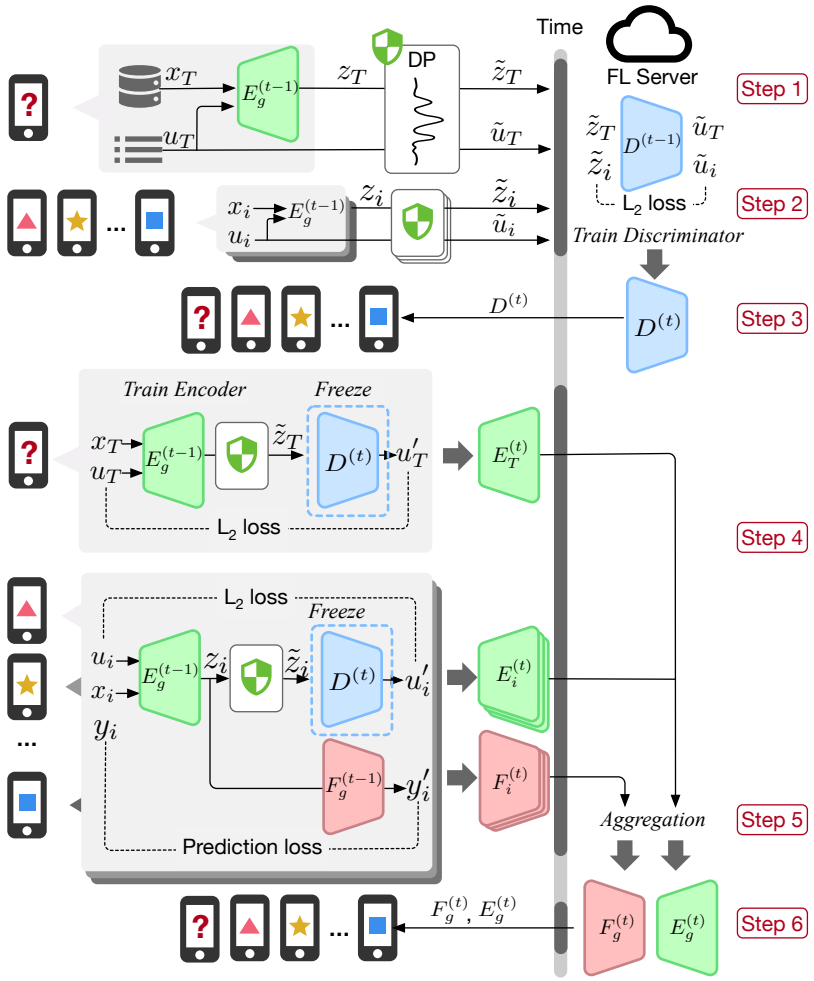

Steps 1–2

Upload, then train the discriminator

Clients upload noised embeddings and indices; the server trains D.

Steps 3–4

Send D back; the target client trains

D is frozen; the unlabeled client updates its encoder to fool it.

Step 4

Source clients: also fit the labels

Encoder and predictor minimize the prediction loss as well.

Steps 5–6

Aggregate and broadcast

The server averages the local models and the next round starts.

Workflow of round t, as in Figure 2 of the paper. The next four slides walk through it.

[1] Q. Yu et al., “FedIndex: Federated Domain Adaptation with Continuous Domain Indices,” *TMLR*, 2026.

Steps 1–2: Private Uploads Train D

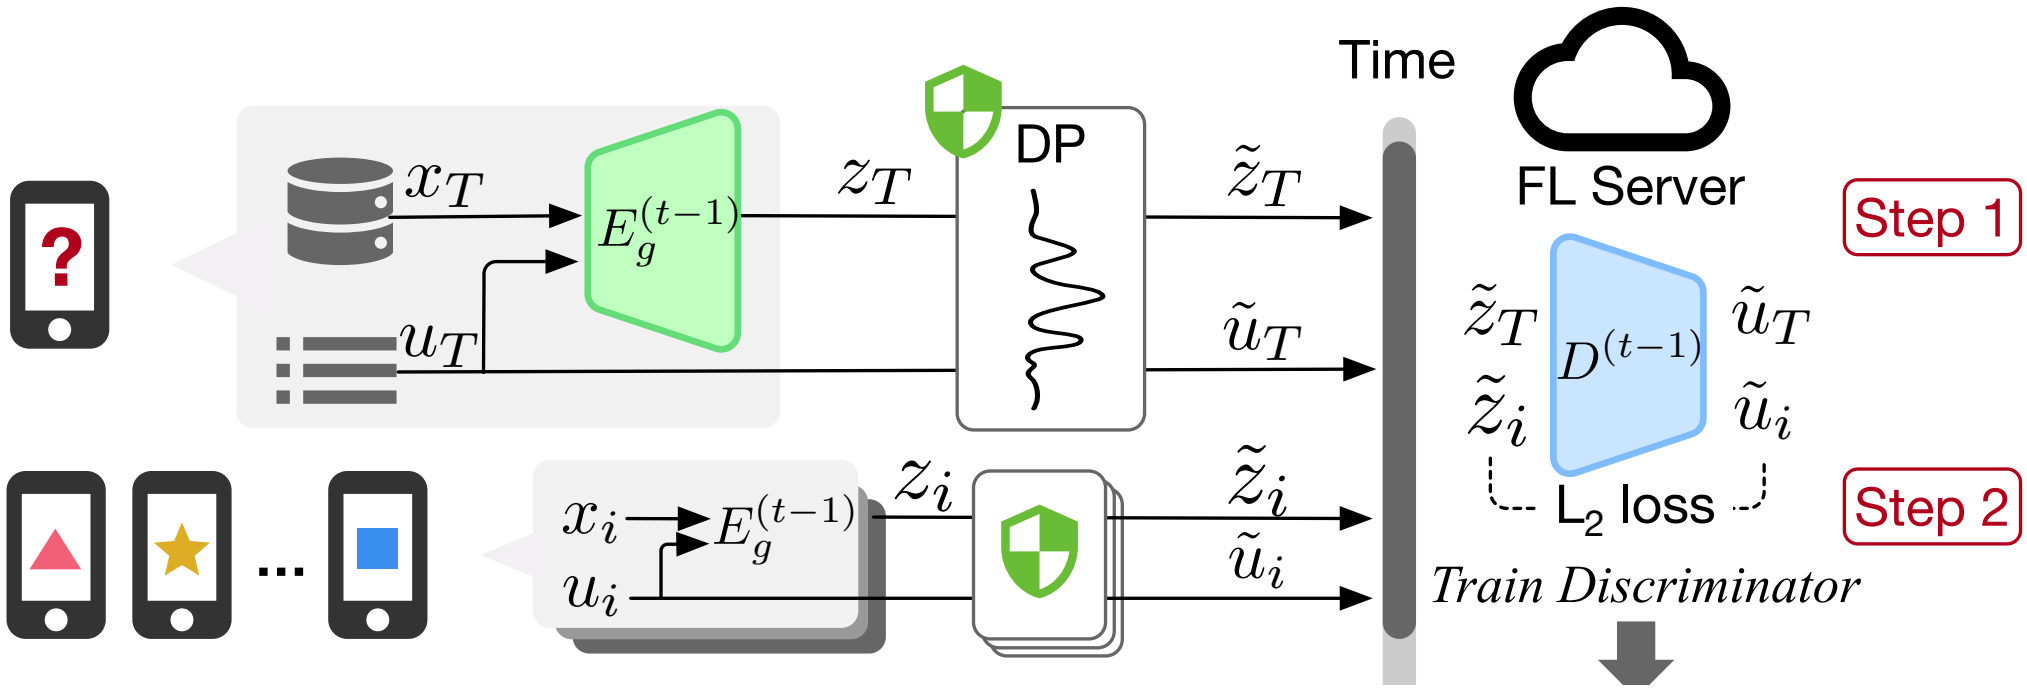

1 Privatize and upload
Every client encodes its data, adds Laplace noise, and uploads $\tilde{z}_i$ and $\tilde{u}_i$.

2 Train the global discriminator
The server regresses $\tilde{u}$ from the noised embeddings of all clients (L2 loss).

[1] Q. Yu et al., "FedIndex: Federated Domain Adaptation with Continuous Domain Indices," *TMLR*, 2026.

Steps 3–4: The Target Client Fools D

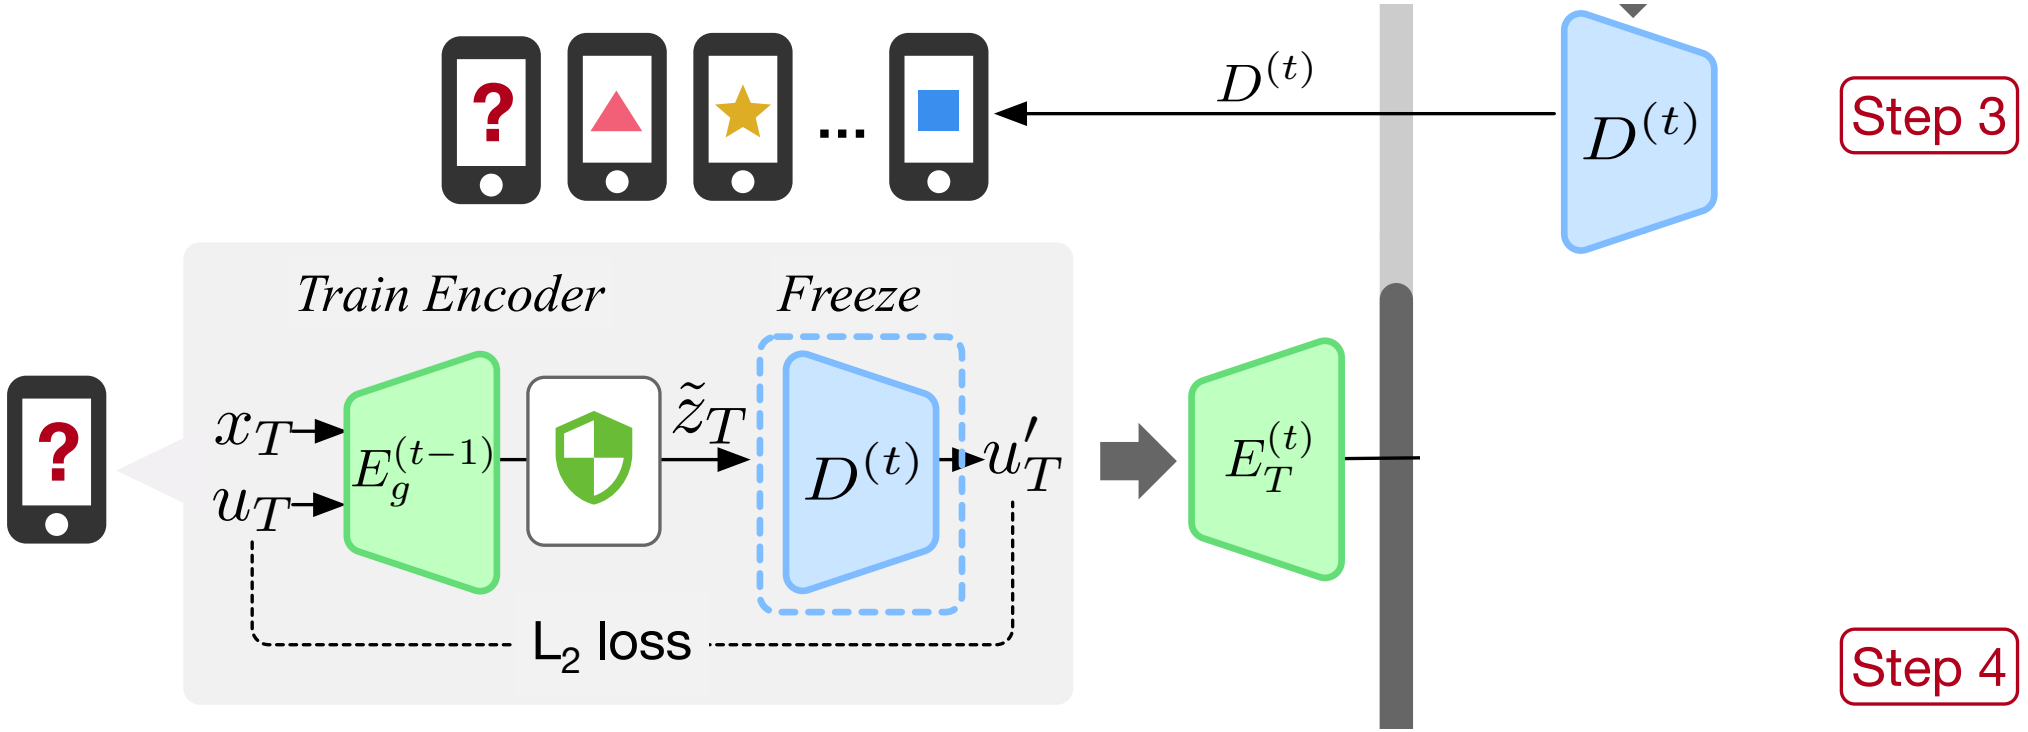

3 **Broadcast**
The server sends the new $D^{(t)}$ to every client.

4 **Train locally: target client**
No labels: only the encoder is trained, to fool the frozen $D^{(t)}$.

[1] Q. Yu et al., "FedIndex: Federated Domain Adaptation with Continuous Domain Indices," *TMLR*, 2026.

Step 4: Source Clients Also Fit Labels

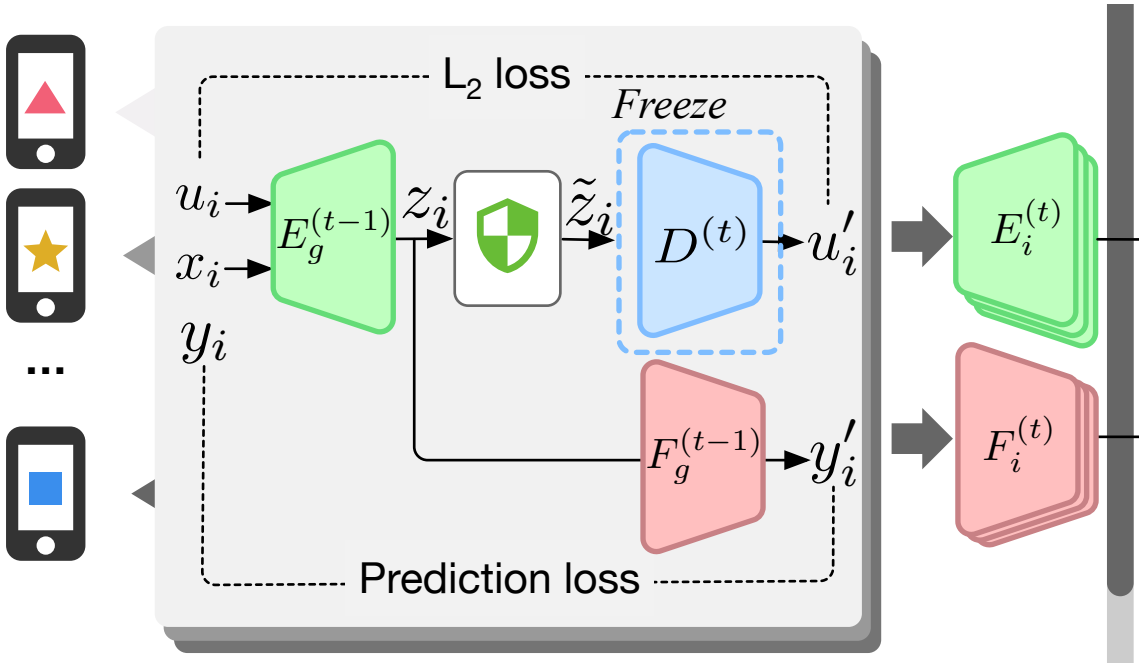

4 **Train locally: source clients**
The encoder fools the frozen $D^{(t)}$; encoder and predictor also fit the labels y_i.

$$\min_{E_i,\,F_i} \underbrace{V_p(E_i, F_i)}_{\text{Fit the labels (source clients)}} - \underbrace{\lambda_d\, V_d(D, E_i)}_{\text{Fool the frozen discriminator}}$$

[1] Q. Yu et al., "FedIndex: Federated Domain Adaptation with Continuous Domain Indices," *TMLR*, 2026.

Steps 5–6: Aggregate and Start Again

5 Aggregate
The server averages the encoders of all clients and the predictors of the source clients into $E_g^{(t)}$ and $F_g^{(t)}$.

6 Broadcast
Every client receives $E_g^{(t)}$, $F_g^{(t)}$ and the next round starts.

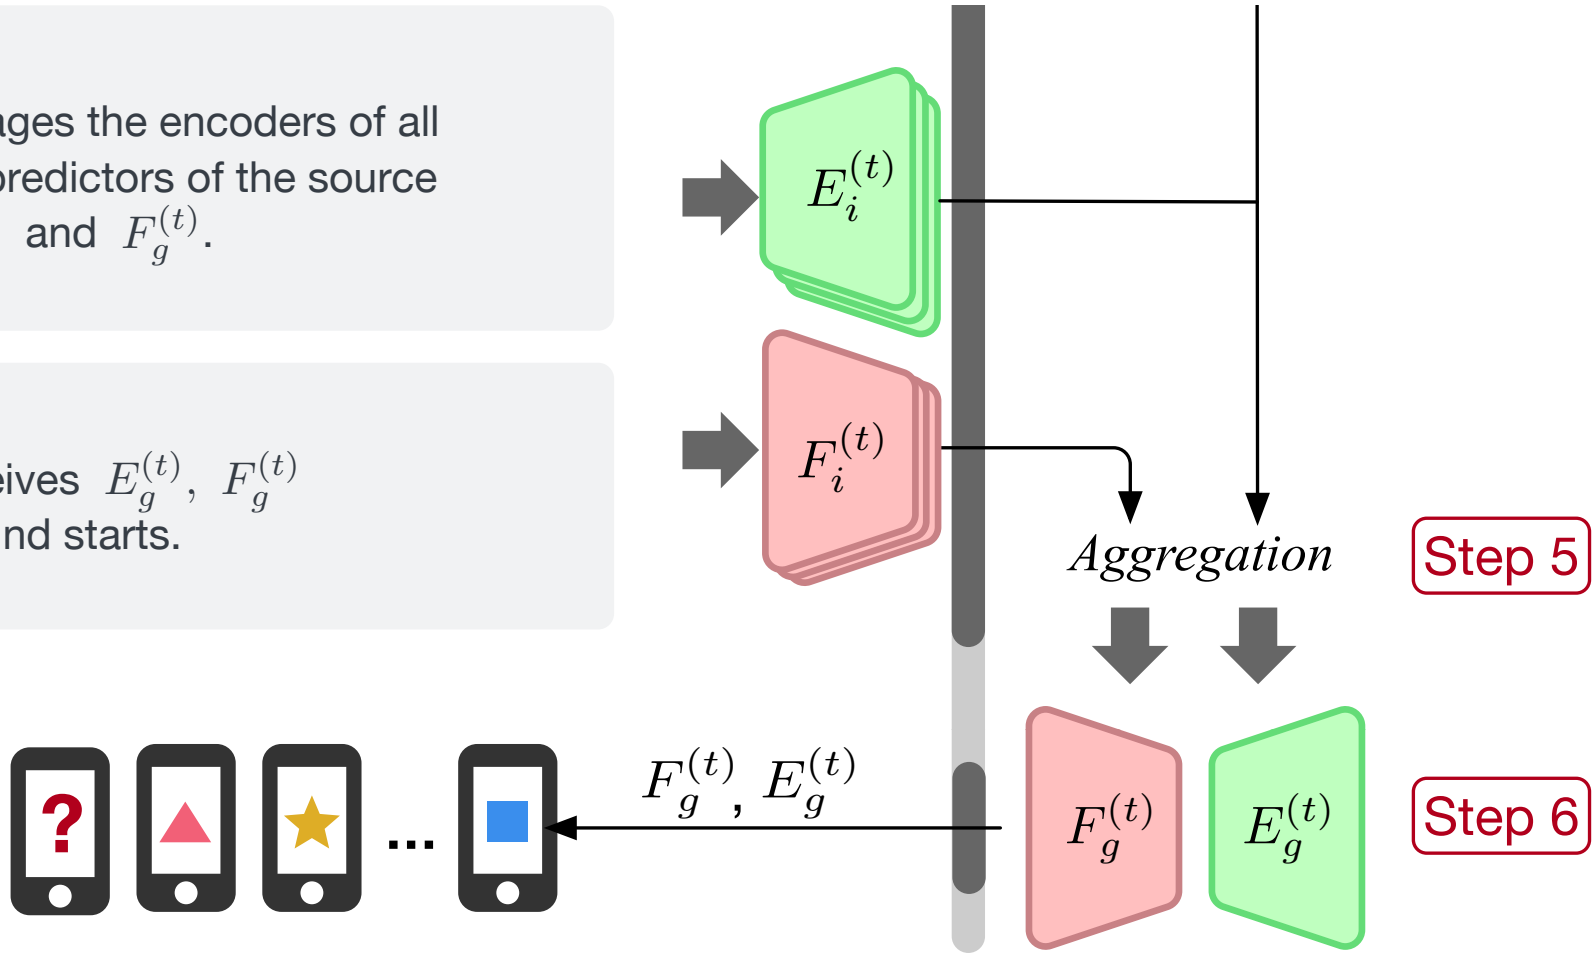

[1] Q. Yu et al., "FedIndex: Federated Domain Adaptation with Continuous Domain Indices," *TMLR*, 2026.

Differential Privacy on Every Upload

Embeddings



$\tau_i \sim \mathrm{Lap}(0,\ \Delta f/\epsilon)$, redrawn every round

Recomputed in every round, so each upload gets fresh noise

Domain indices



$\tau_u \sim \mathrm{Lap}(0,\ \Delta f/\epsilon)$, drawn once and reused

Static across rounds, so the same noised index is uploaded each time

One Laplace Mechanism Protects Both Uploads

Kolmogorov–Smirnov tests give $p < 0.05$ for every ε from 0.01 to 1. DP guards against distribution- and membership-inference attacks.

[1] Q. Yu et al., "FedIndex: Federated Domain Adaptation with Continuous Domain Indices," *TMLR*, 2026.

FedIndex Aligns Domains at Equilibrium

Optimal alignment

At the global optimum, embeddings are mean-aligned across all clients

$$\mathbb{E}[u \mid \tilde{z}] = \mathbb{E}[u], \quad \forall\, \tilde{z}$$

Robust to DP noise

Zero-mean Laplace noise only adds a constant to the objective

$$C_d(E) = \mathbb{E}_{\tilde{z}}\, \mathbb{V}[u \mid \tilde{z}] + \mathbb{V}[\tau_u]$$

Aligned Domains, No Label Information Lost

With the index independent of the label, the optimal encoder keeps all label information

Left: Theorem 5.2, for zero-mean noise. Right: Corollary 5.2; the constant does not affect encoder gradients. Below: Theorem 5.3.

[1] Q. Yu et al., "FedIndex: Federated Domain Adaptation with Continuous Domain Indices," *TMLR*, 2026.

Best Federated Accuracy on All Six Datasets

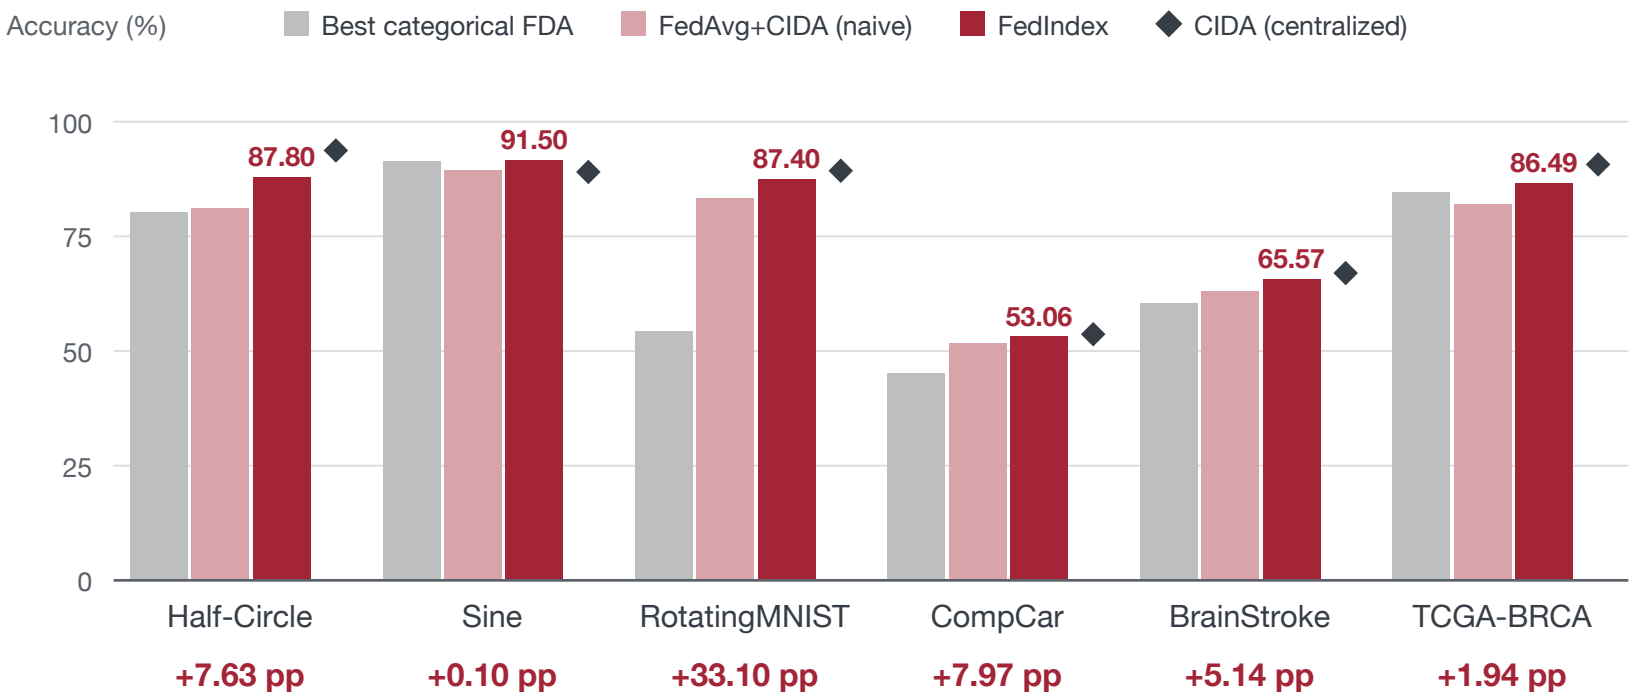

87.40%

RotatingMNIST accuracy

+33.10

points over the best categorical baseline

Feature shift. Best categorical FDA baseline per dataset among FedAvg+MDD, FedAvg+DANN, FADE+DANN, FedDA and FedGP.

[1] Q. Yu et al., "FedIndex: Federated Domain Adaptation with Continuous Domain Indices," *TMLR*, 2026.

Larger Gains Under Joint Distribution Shift

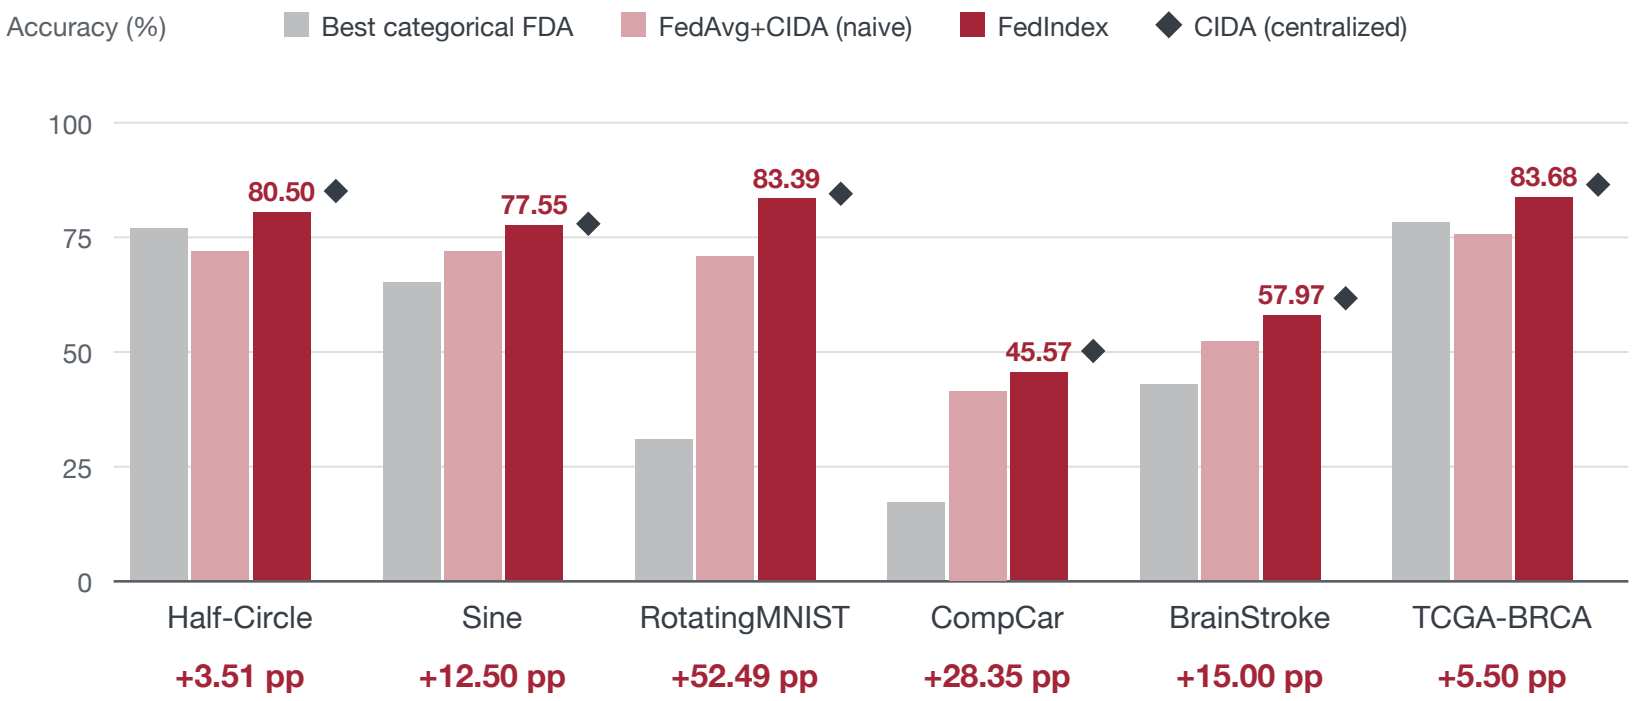

83.39%

RotatingMNIST accuracy

+52.49

points over the best categorical baseline

Joint distribution shift adds label shift on top of feature shift. Best categorical FDA baseline chosen per dataset, as before.

[1] Q. Yu et al., "FedIndex: Federated Domain Adaptation with Continuous Domain Indices," *TMLR*, 2026.

The Continuous Index Drives the Gains

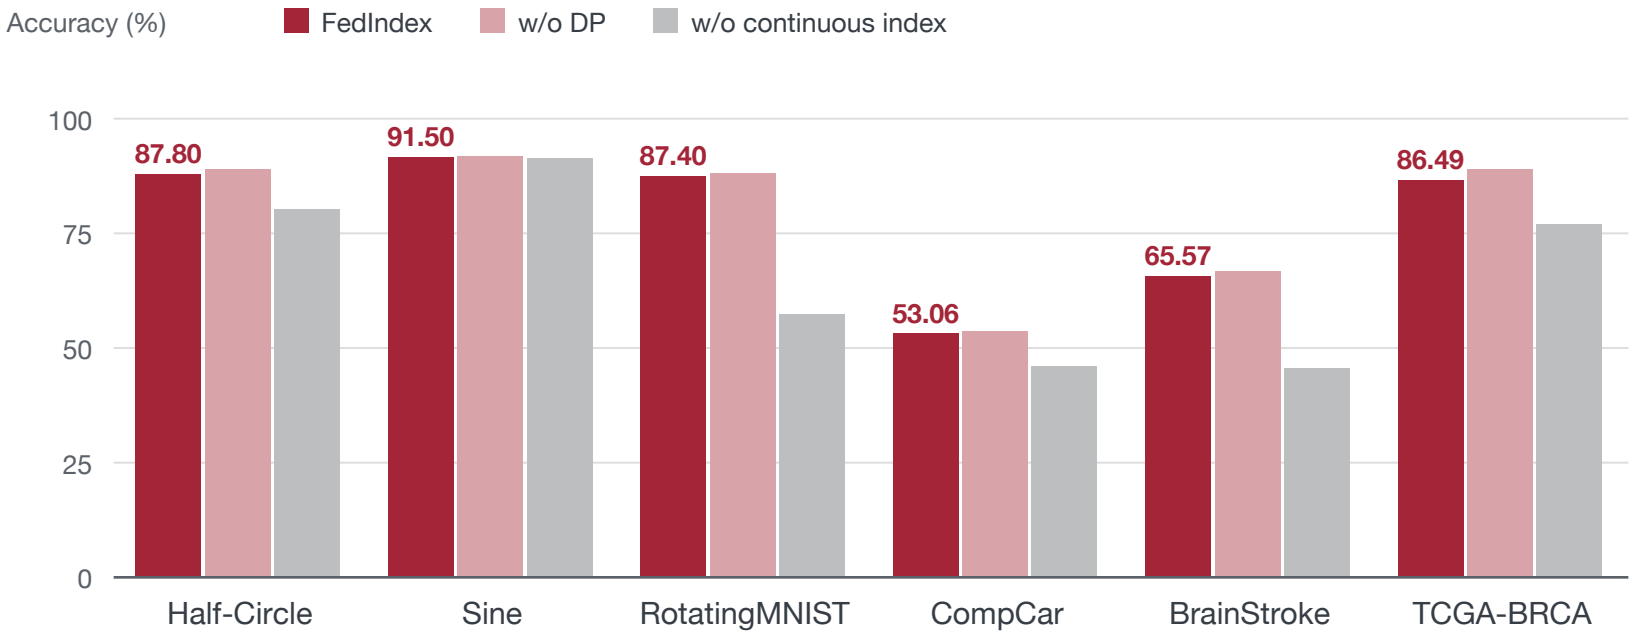

−30.00

pp without the continuous index

−2.48

pp at most from DP noise (ε = 1)

Largest drop without the index: RotatingMNIST. Largest cost of DP: TCGA-BRCA. Feature shift, six datasets.

[1] Q. Yu et al., “FedIndex: Federated Domain Adaptation with Continuous Domain Indices,” *TMLR*, 2026.

FedIndex

Continuous indices, one global discriminator

FedIndex aligns continuously indexed domains across clients, with differential privacy on every upload.

+3.5 to +52.5 pp accuracy

over the best categorical FDA baseline under joint shift

qyu13@stevens.edu · github.com/IntelliSys-Lab/FedIndex

[1] Q. Yu et al., “FedIndex: Federated Domain Adaptation with Continuous Domain Indices,” *TMLR*, 2026.



TMLR
Transactions on Machine Learning Research
Published June 2026

STEVENS
INSTITUTE OF TECHNOLOGY
1870

Thank You

Code
github.com/IntelliSys-Lab/FedIndex

Homepage
ryougish1k1.github.io

IntelliSys Lab
intellisys.haow.us

Stevens Institute of Technology
1 Castle Point Terrace, Hoboken, NJ 07030