

EFU: Enforcing Federated Unlearning via Functional Encryption

Samaneh Mohammadi · **Vasileios Tsouvalas** · Iraklis Symeonidis · Ali Balador · Tanir Ozcelebi · Francesco Flammini · Nirvana Meratnia

RISE Research Institutes of Sweden | Mälardalen University | Eindhoven University of Technology

ACM CIKM 2025

01 Motivation

The problem

A client’s unlearning update is ineffective if omitted from aggregation. *EFU* addresses this protocol gap while studying how forgetting affects retained utility and observable model changes.

Contributions

- **Aggregate authorization:** Bind encrypted updates to an agreed aggregation function.
- **Local unlearning:** Combine forget loss, adversarial training and inverted Memory Aware Synapses (MAS) regularization.
- **Forgetting evaluation:** Compare sample-, class- and task-level removal with full retraining.

Threat model

The server is honest-but-curious (HbC) and may inspect updates or try to omit an unlearning contribution. Clients are trusted to perform key generation, encryption and unlearning correctly; no trusted third-party authority (TPA) is assumed.

Protection boundary

Functional keys authorize one round’s aggregate. Update-type hiding assumes indistinguishable outputs and excludes full collusion; visible mappings and side channels remain separate concerns. This does not compel deployment or certify complete data erasure.

Scope


Results average three runs without reporting dispersion. Similar retained/forgotten accuracies do not establish statistical equivalence to retraining.

Notation

Dirichlet distribution with concentration a (Dir(a)) · Labeled Faces in the Wild (LFW) · functional encryption (FE) · federated unlearning (FU) · vision transformer (ViT) · unlearning loss (UN) · cross-entropy (CE)

02 EFU

Bind updates to the authorized aggregate



Technical core

$$\mathcal{L}_{UN} = -\mathcal{L}_{CE}(D^u) + \mathcal{L}_{CE}(D^{u,adv}) + \sum_j \bar{\Omega}_j(\theta_j - \tilde{\theta}_j)^2$$

D^u : forget data; $D^{u,adv}$: adversarial variants; $\tilde{\theta}$: starting weights; $\bar{\Omega}$: inverted MAS importance. Task forgetting also optimizes retained-task loss.

$$\theta_j^{r+1} = \text{Dec}(\{\theta_{i,enc}^r[j]\}_i, dk_f)$$

Encrypted client updates and the function key dk_f yield the agreed aggregate. The binding concerns decryption, not subsequent deployment.



Paper



Code

03 Evaluation

95.37%

Retained gender accuracy

LFW · task forgetting

5.12%

Forgotten identity accuracy

Same *EFU* model

Class forgetting with retained utility

Method	Retain accuracy (%)	Forget accuracy (%)
Full retrain	94.82	13.72
PGD	91.82	14.31
FedOSD	93.06	12.77
EFU + FedOSD (Ours)	92.73	12.54
EFU (Ours)	92.43	13.06

CIFAR-10 · ResNet-18 · forget 10% of classes · 10 clients (1 unlearning) · 20% participation · Dir(0.1), 100 rounds; 10 unlearning rounds from round 50. Means of 3 runs; higher retain / lower forget is better.

Forget identity, retain gender

Method	Gender accuracy (%)	Identity accuracy (%)
Full retrain	96.86	4.71
PGD	94.81	10.32
FedOSD	95.02	5.01
EFU (Ours)	95.37	5.12

LFW · ViT-B/32 · 30 clients, all unlearning identity · full participation. Ten unlearning rounds from round 50; means of 3 runs.

Takeaway

EFU retains **95.37% gender accuracy** while identity accuracy falls to **5.12%** in LFW task forgetting. Its **authorized encrypted aggregation** incorporates the agreed updates, while complete erasure and subsequent model deployment remain separate requirements.