

AutoEncoder-Compressed Parallel Split Learning for Pre-trained Model Fine-Tuning

Bas Meuwissen · Vasileios Tsouvalas · Nirvana Meratnia

Eindhoven University of Technology

4th Workshop on Advancements in Federated Learning (WAFL), ECML-PKDD 2026

01 Motivation

The problem

Split learning moves expensive computation to a server, but every training step still transfers activations and gradients. Compression must preserve the feature distribution learned during pretraining.

Contributions

- **Bidirectional compression:** reduce activation and gradient traffic through an aligned autoencoder (AE).
- **Client adaptation:** retain local encoders while sharing one averaged server decoder.
- **Deployment evaluation:** test global and local inference across four vision benchmarks.

Pretrained features are fragile

A randomly initialized AE can destroy useful representations. The compressor must match the pretrained feature distribution before downstream fine-tuning.

Two deployment settings

Global inference removes the AE; local inference retains it. The first three benchmark datasets use global evaluation, while FEMNIST uses local evaluation.

Scope


Evaluated on supervised vision tasks with vision transformer (ViT)-B/32; multimodal and language adaptation remain untested. Public-data alignment and client warm-up are required.

Notation

parallel split learning (PSL) · mean squared error (MSE) · standard deviation (SD)

02 AE-PSL

Align the compressor before fine-tuning



Technical core

$$\mathcal{L}_{\text{rec}} = \frac{1}{|\mathcal{D}|} \sum_{x \in \mathcal{D}} \|H(x) - D_{\psi}(E_{\phi}(H(x)))\|_2^2$$

H: frozen pretrained features; E/D: encoder/decoder. Minimize reconstruction error on public data, then each client's data.



Paper



Code

03 Evaluation

≈10×

communication reduction

Relative to uncompressed split fine-tuning.

82.7 ± 0.2%

CIFAR-100 accuracy

5 clients; matches the uncompressed reference.

Accuracy at approximately 10× communication reduction

Method	CIFAR-100	Food101	SUN397	FEMNIST
No compression (R=1)	82.7 ± 0.2	69.1 ± 0.9	67.0 ± 0.6	87.9 ± 0.3
C3-SL	50.9 ± 1.2	40.3 ± 1.2	41.3 ± 2.1	83.5 ± 0.1
Rand-Top-K	76.2 ± 0.4	60.2 ± 0.2	62.9 ± 1.0	87.2 ± 0.4
ADC	68.4 ± 0.4	54.6 ± 1.6	52.9 ± 1.1	81.1 ± 0.9
AE-PSL (Ours)	82.7 ± 0.2	69.6 ± 1.2	66.7 ± 0.3	88.0 ± 0.1

Top-1 accuracy (%), mean ± SD, 3 seeds; ViT-B/32, 5 clients. Compressed methods R≈10. Global: first three datasets, AE bypassed; local: FEMNIST, AE active.

Alignment prevents feature mismatch

Ablation	CIFAR-100	FEMNIST
Unaligned AE	1.4 ± 0.3	86.9 ± 2.0
+ General alignment	75.2 ± 0.4	87.1 ± 1.1
+ Freeze AE	79.6 ± 0.9	88.0 ± 0.6
+ Client alignment (Ours)	82.2 ± 0.7	88.0 ± 0.7

Cumulative ablation; accuracy (%), reported ±, 3 seeds; ViT-B/32, 5 clients, R=15. CIFAR-100: global; FEMNIST: local.

Takeaway

AE-PSL matches uncompressed CIFAR-100 accuracy at **≈10× lower communication**. **Alignment, freezing and client adaptation** each improve the reported CIFAR-100 ablation.