NEURIPS 2026

Focus on Where You Aggregate:

Restricted SAM for Non-IID Federated Learning

Haotong Wen, Junkang Liu, Yi Xu, Xiao Liu, Longkun Guo, Kewen Liao

Presented by Haotong Wen

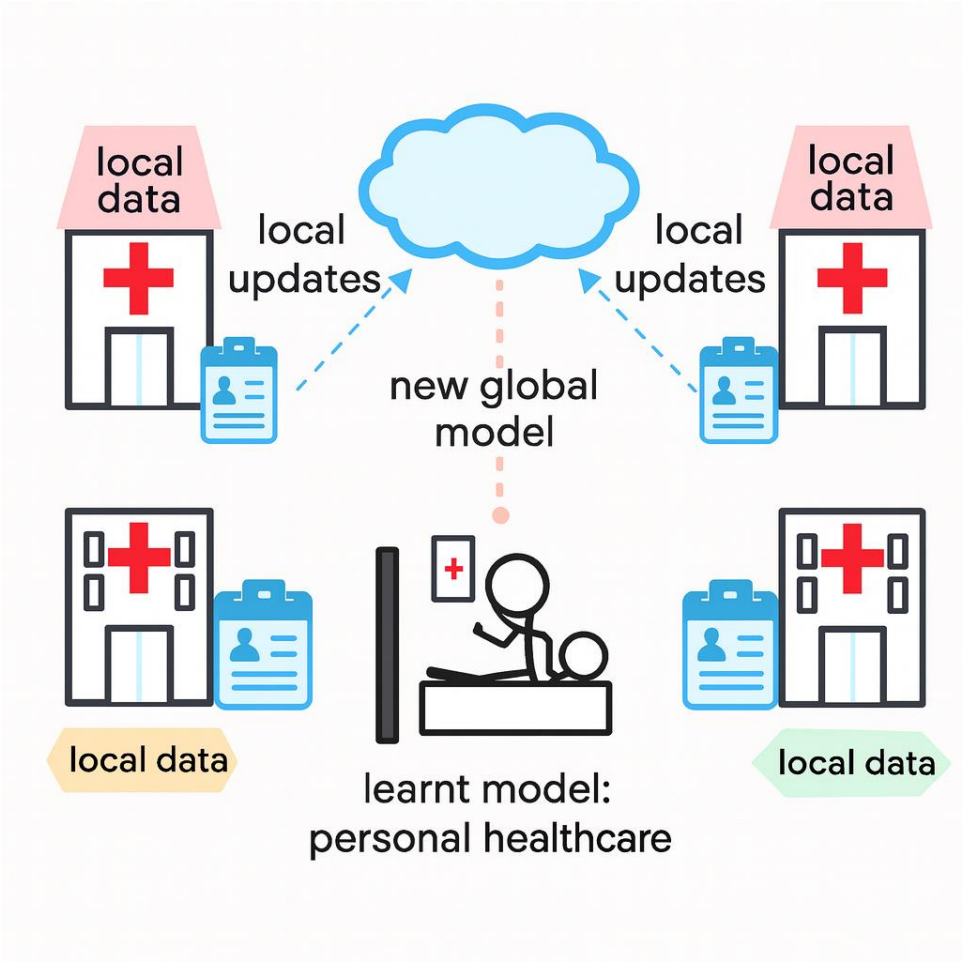
Anhui University

Deakin University

Tianjin University

Fuzhou University

Federated learning with heterogeneous clients



Local data, shared model

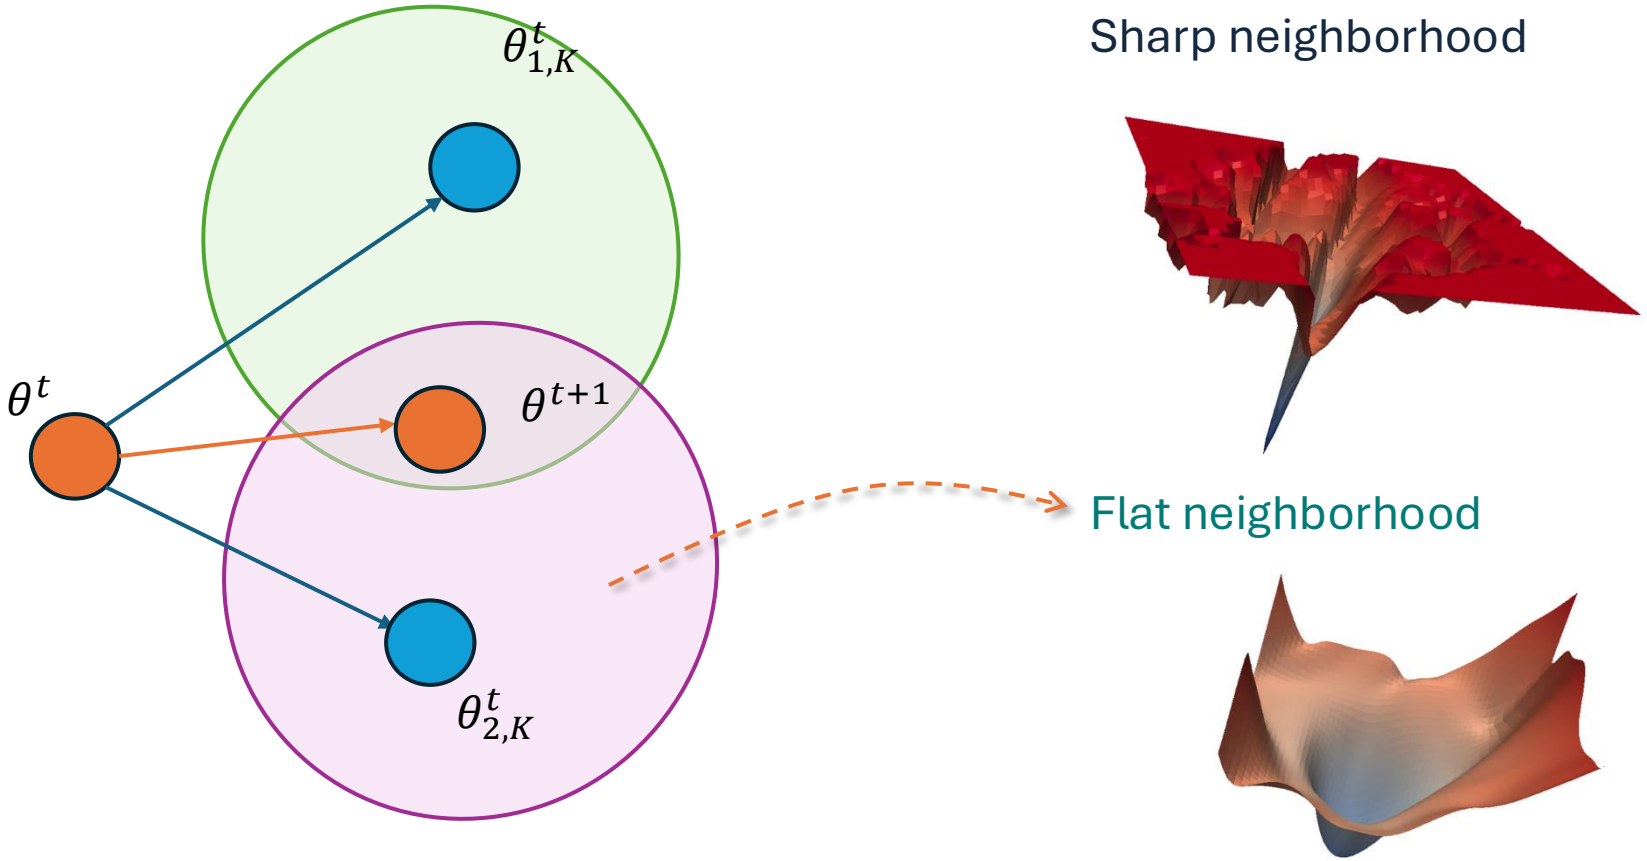
Clients train on their own data.

The server aggregates local models.

Different data can pull clients apart.

Federated setup. Illustration from the supplied presentation.

FedSAM: local flatness before aggregation

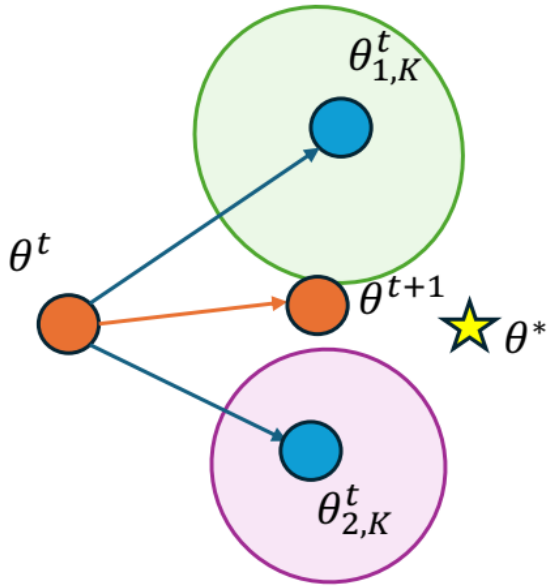


Overlapping local neighborhoods can cover the aggregate.

SAM / Foret et al. (2021). FedSAM / Qu et al. (2022). Schematic illustration.

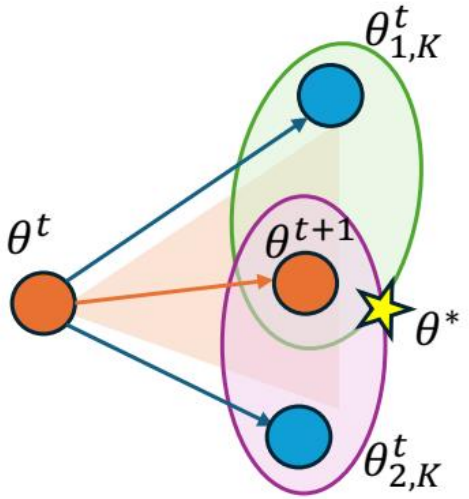
Local flatness can miss the aggregation region

FedSAM



Search the full local neighborhood

FedRSAM

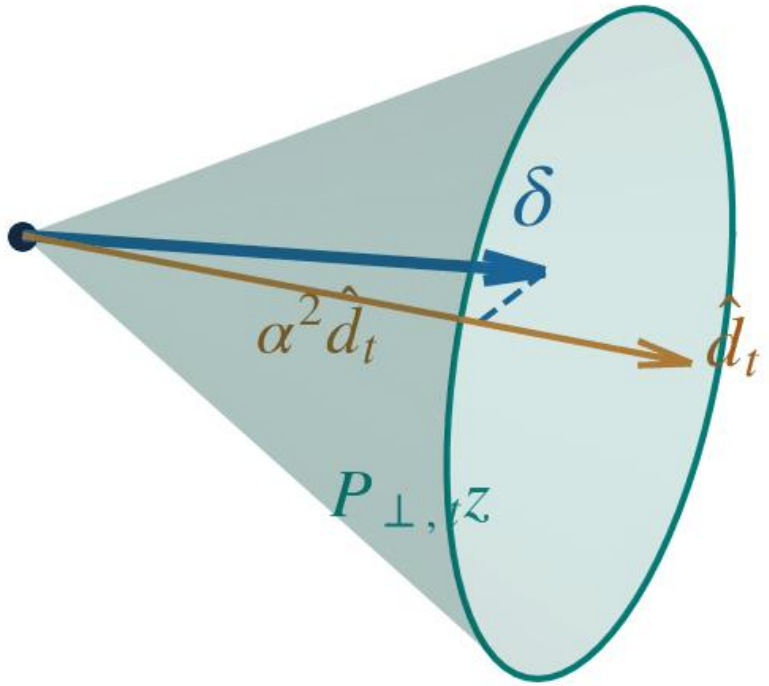


Search toward likely aggregation

Clients train locally. The server combines their models.

Paper Figure 1(a–b). Blue circles mark client models. Orange circles mark the server model.

Restricted perturbations inside the cone



Model perturbation

$$\theta \longmapsto \theta + \delta(\alpha, z)$$

Cone parameterization

$$\delta(\alpha, z) = \alpha^2 \hat{d}_t + P_{\perp,t} z$$

One ascent step per minibatch

$$\alpha^{(1)} = \alpha^{(0)} + \eta_\alpha \frac{\partial L_{\text{in}}}{\partial \alpha}$$

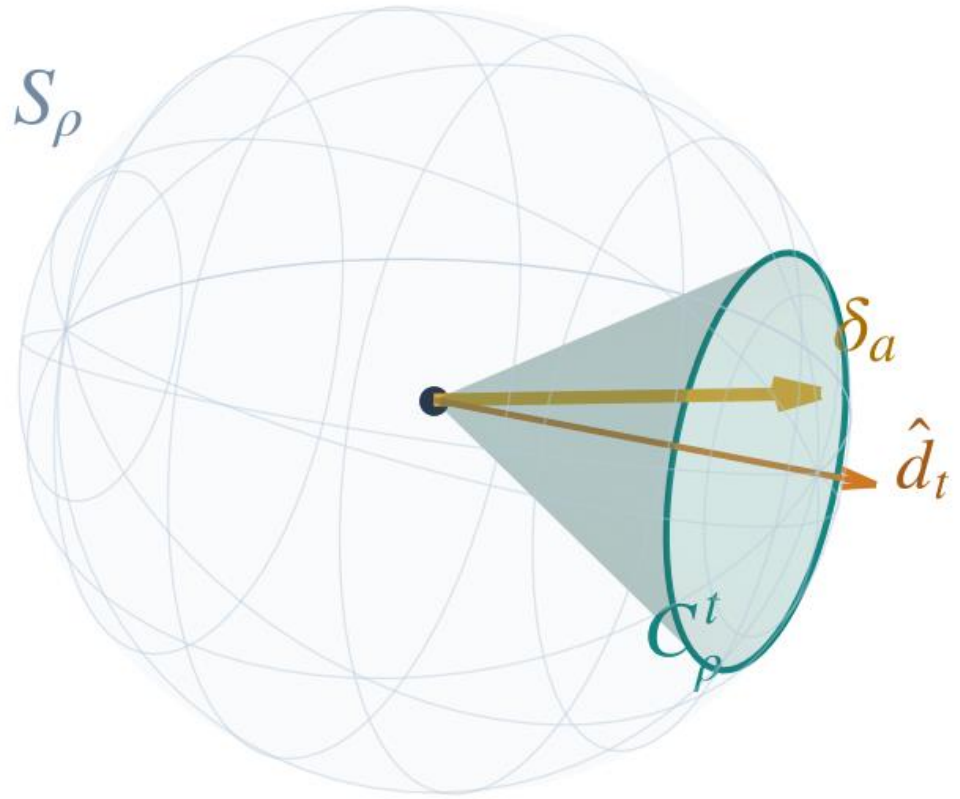
$$z^{(1)} = \Pi_{[-\sqrt{v_t},\, \sqrt{v_t}]}\left(z^{(0)} + \eta_z \frac{\partial L_{\text{in}}}{\partial z}\right)$$

Then cap the perturbation at radius ρ.

Update the model at the perturbed point.

Schematic geometry. The actual residual constraint is coordinate-wise. Paper Eqs. (7)–(14).

Why restrict the perturbation directions?



The ideal perturbation δ_a

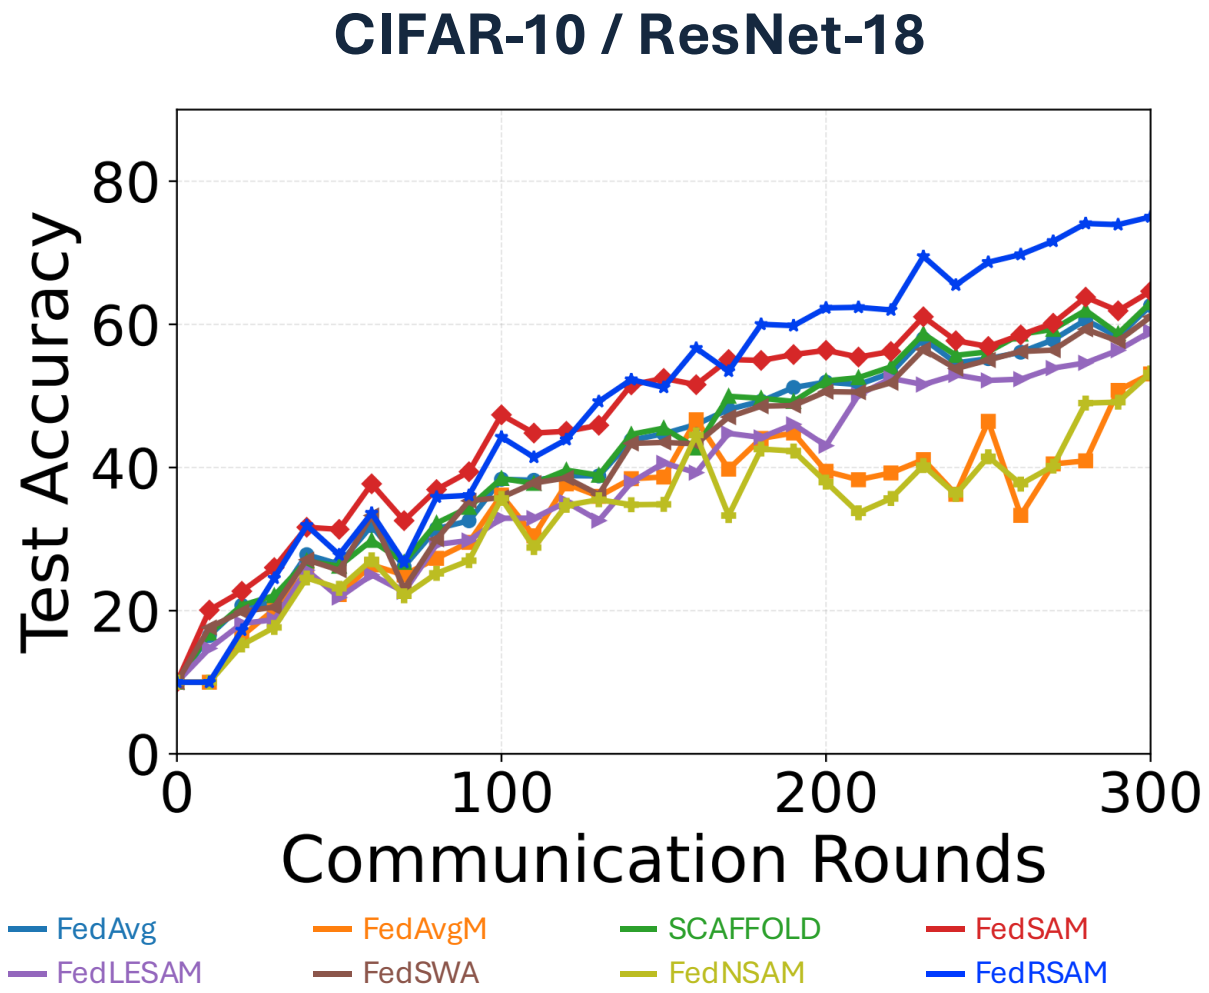
points along the realized aggregation displacement.

The restricted region

has a tighter expected aggregation-sensitivity error bound under the stated assumptions.

Paper Theorem 1 and Proposition 1. Uniform directional references, positive axis alignment and a symmetric residual section.

Vision benchmark accuracy



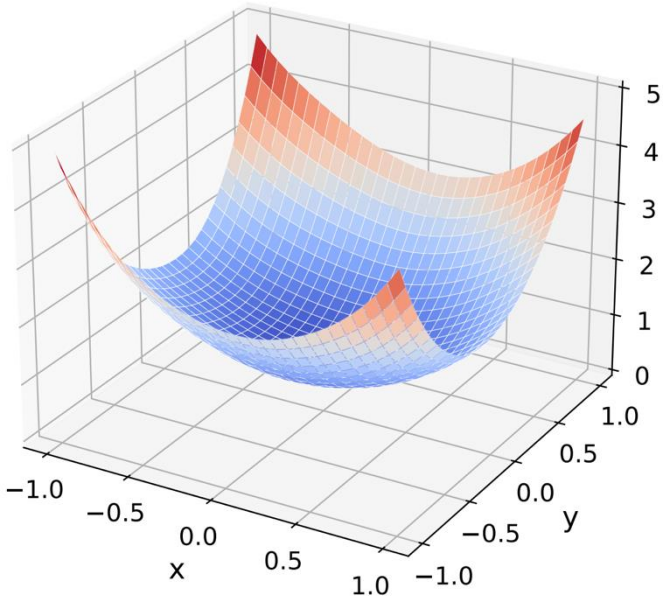
Test accuracy (%)

Dataset	FedSAM	FedRSAM
CIFAR-10	64.62	**75.01**
CIFAR-100	37.90	**52.51**
Tiny-ImageNet	74.35	**81.55**

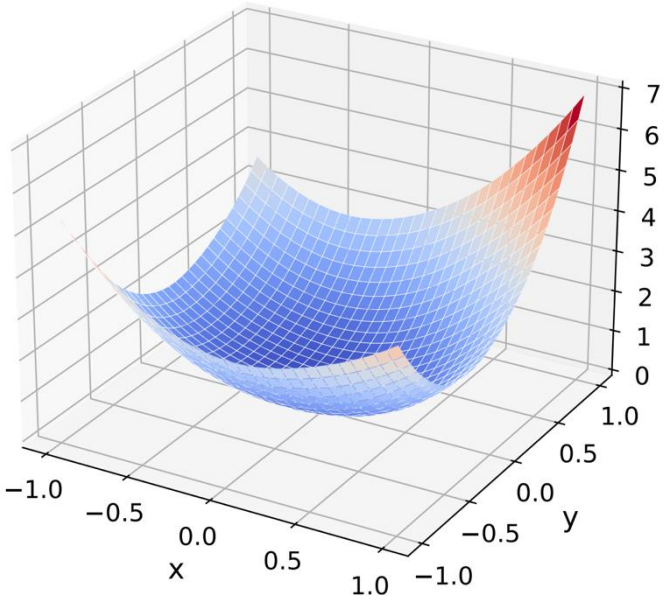
Paper Fig. 2(b) and Table 1. CIFAR uses ResNet-18; Tiny-ImageNet uses ViT-B.

Loss geometry and aggregation behavior

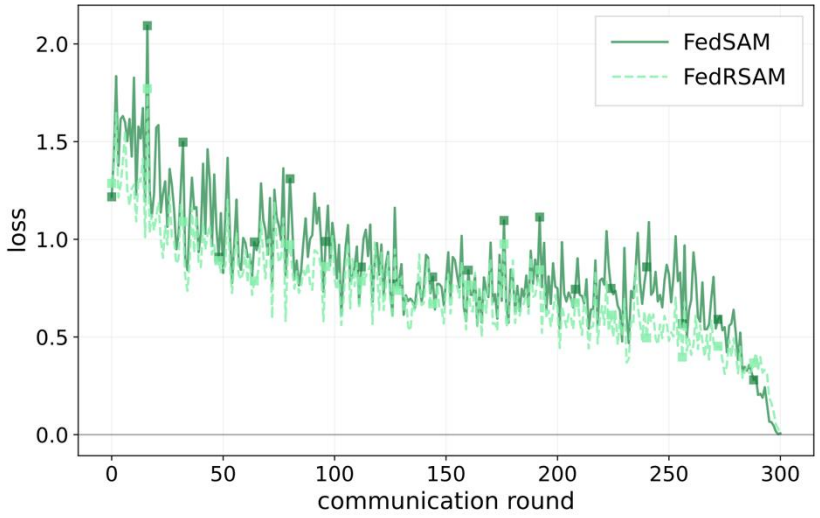
FedSAM



FedRSAM



Aggregation loss gap



FedSAM flattens a broader neighborhood.

FedRSAM generally has a smaller aggregation gap.

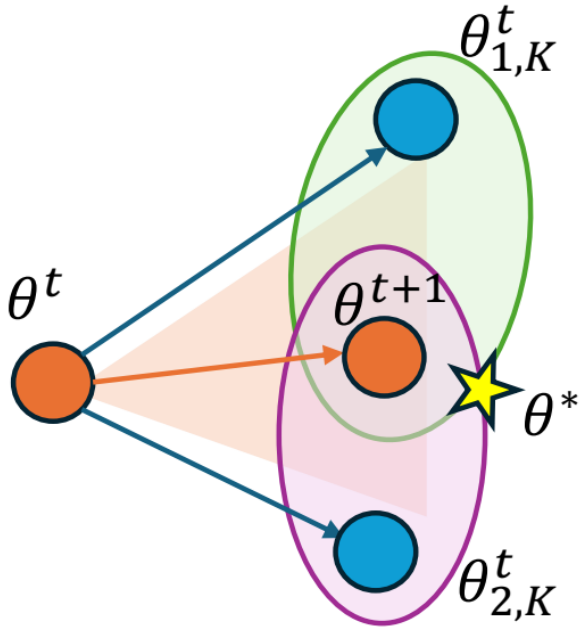
Paper Fig. 3(a–c), CIFAR-10. The curve shows the aggregation loss gap discussed in §5.3.

Conclusions and next steps

Conclusions

Target the region the server is likely to visit.

Combine a historical axis with
variance-constrained residual exploration.



Looking ahead

Improve the geometry estimate
under very sparse participation.

github.com/Doctorwngd/FedRSAM-official

Thanks for listening

FedRSAM

Haotong Wen