Federated Neural Radiance Fields

Lachlan Holden, Feras Dayoub, David Harvey, and Tat-Jun Chin

AUSTRALIAN INSTITUTE FOR MACHINE LEARNING

THE UNIVERSITY of ADELAIDE

Problem

Image filesizes are large.

Can NeRFs be trained distributed across nodes in a bandwidth-efficient way?

Yes. We contribute **FedNeRF**, a novel federated learning algorithm for NeRF.

Centralised vs federated learning

In **centralised learning**, data-collecting edge nodes transmit data back to a central server for training.

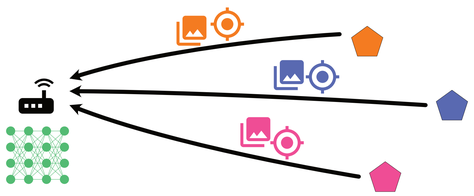

In **federated learning**, edge nodes train on-board and transmit network weights back to a central server, which aggregates them.

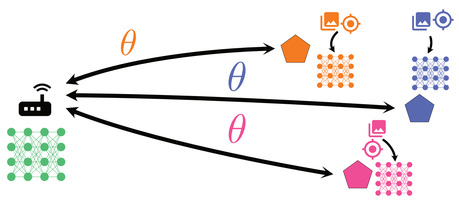

Motivating scenario

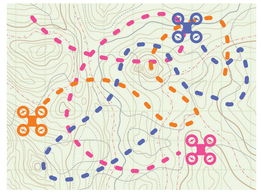

Distributed mapping campaign of a remote area.

Train a single, shared NeRF to represent the scene.

Drones have GPS, take photos, and train on-board.

Take-home message

NeRFs can be trained using federated learning.

This can reduce bandwidth use and training iterations.

FedNeRF algorithm

Initial weights are trained on the server.

Singlular value decomposition discards and freezes weights, for update compression. ▶❶

The frozen weights are transmitted to each client.

Each client takes new photos, then in a loop:

- Each client receives the current weights from the server
- Each client trains their NeRF some more
- Each client transmits their NeRF weights to the server
- The NeRF weights are combined using federated averaging on the server ▶❷

❶ Update compression via SVD

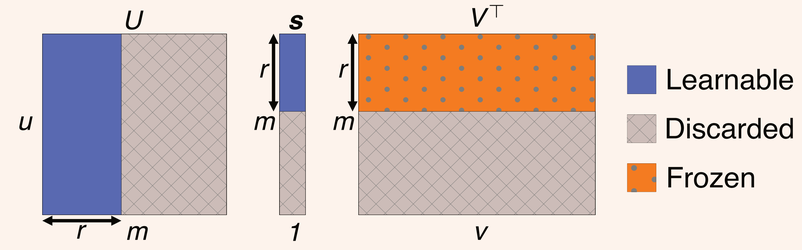

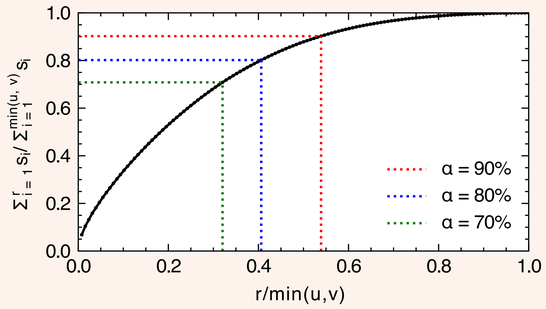

Adaptive rank selection ensures sufficient variance is retained.

❷ Federated averaging

A dataset-size weighted average of learnable weights is taken, as in standard Fed. Avg.

Experimental results

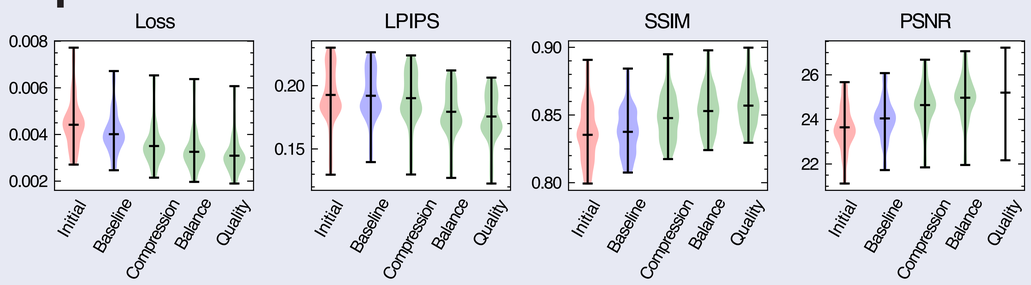

Quality metrics for validation images of the lego dataset, after 20 000 initial + 20 000 further training iterations.