

# Can Fair Federated Learning reduce the need for Personalisation?

Alex Iacob, Pedro P. B. Gusmão, Nicholas D. Lane  
 {aai30,pp524 ndl32}@cam.ac.uk

**TL;DR**

- Federated Learning (FL) does not perform **fairly** across clients
- Local models may have higher accuracy than federated ones
- Fair** and/or **Personalised** Federated Learning address this
- Our contribution:
  - We show Fair FL (FFL) does **not** benefit personalisation
  - We propose using regularisers that **anticipate** personalisation
  - We show **Personalisation-aware** FL (PaFL) outperforms FFL

**Experimental Setup**

- Two datasets: **Reddit**, **FEMNIST**
- Train centralised models using:
  - Q-FedAvg or TERM
  - PaFL: FedAvg+EWC/KD after the halfway round of training
- Personalise with KD, EWC, FB, or None
- Eval **delta** between FL/personalised model and fully local models

**Fair FL**

Fair FL fails to reduce underperforming clients and harms accuracy.

**Reddit**

- Slight **decrease** in average accuracy

**FEMNIST:**

- Doubles** underperforming clients
- Despite **improving** average accuracy

**Personalisation-aware FL**

PaFL either does not harm the number of underperforming clients or halves it.

**Reddit**

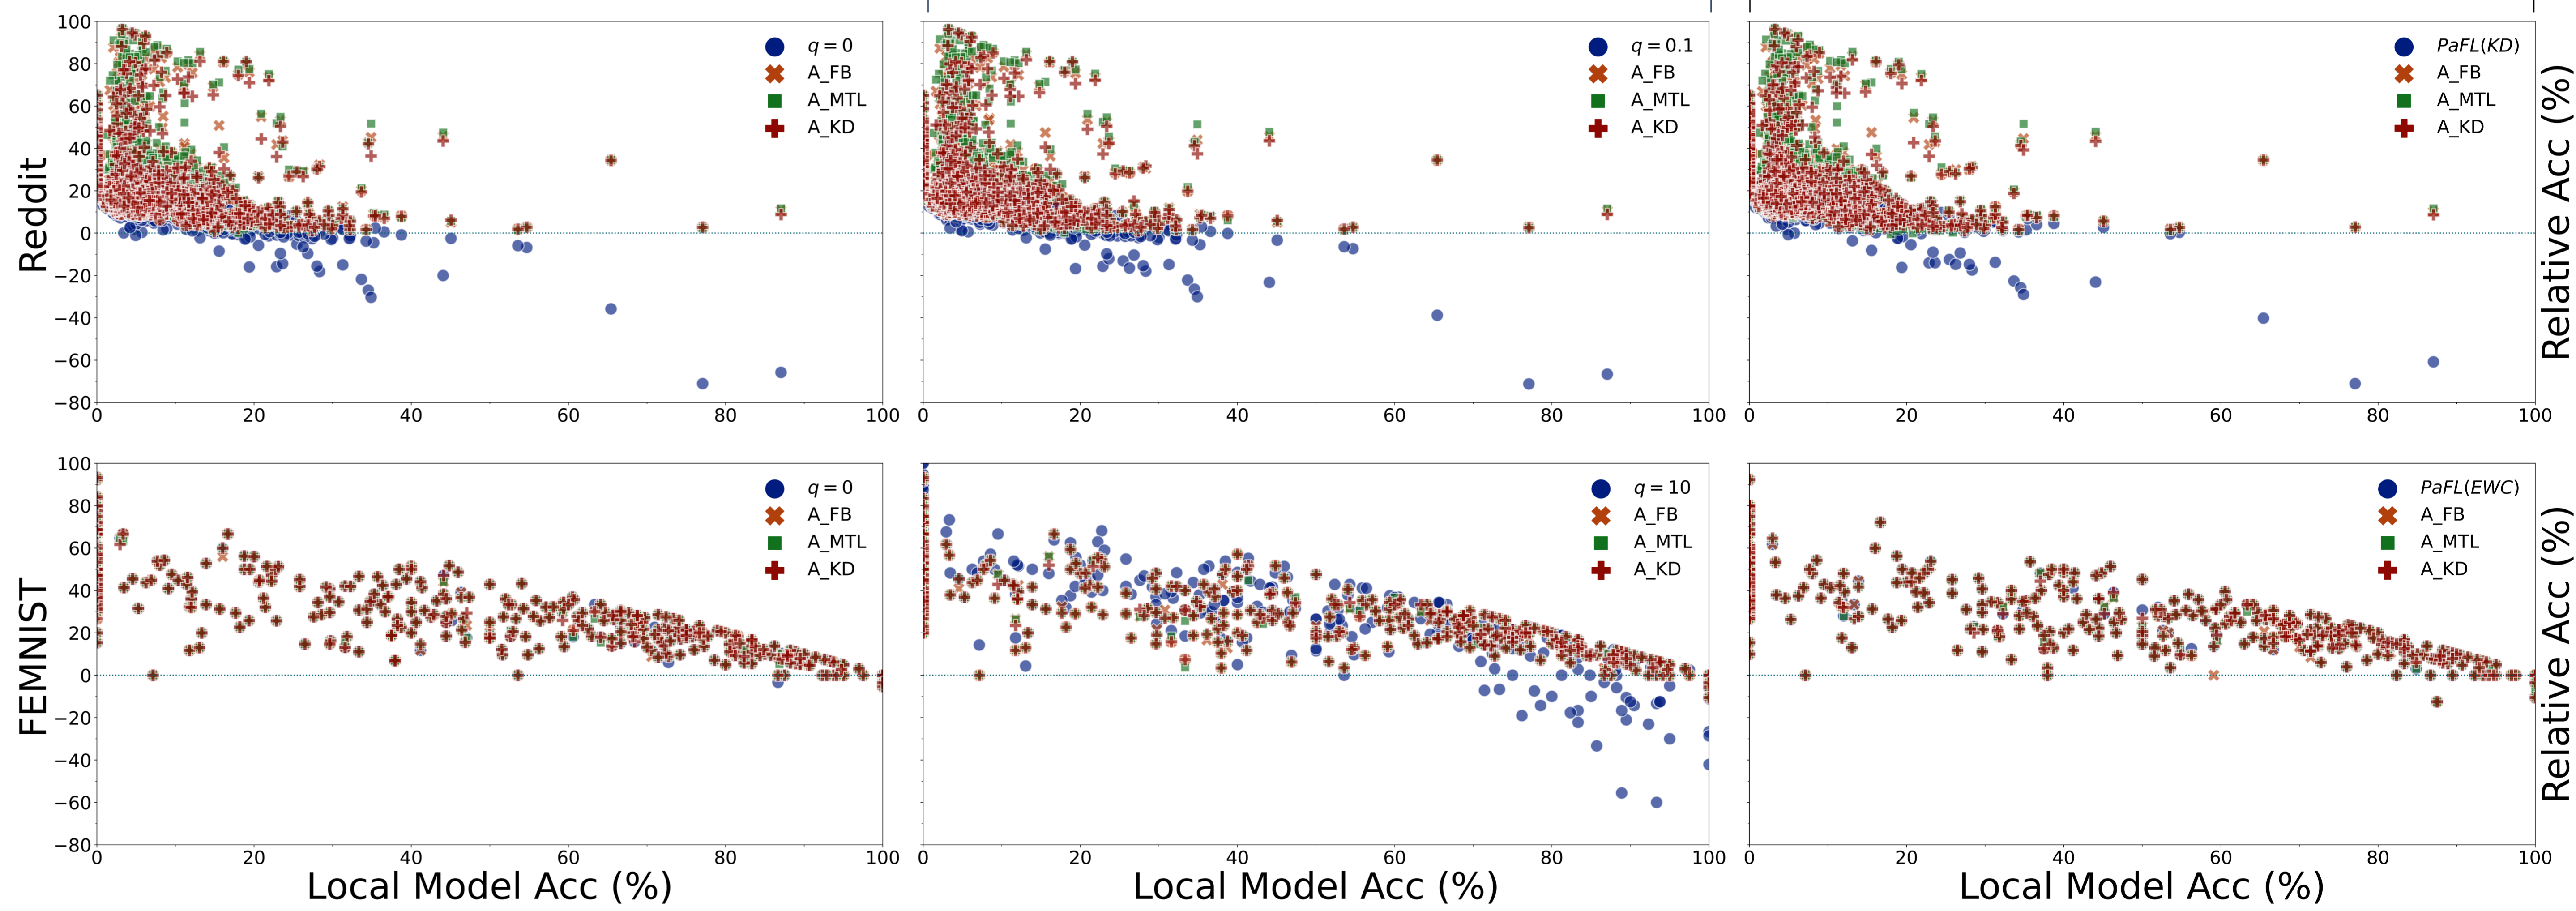
- Halves** underperforming clients

**FEMNIST:**

- Achieves similar average accuracy
- Improves personalisation with EWC

**Fair FL and Personalisation**

- FL trains models directly on client devices and then averages them
- The data distribution of clients in FL is **Non-IID**
- FL models perform worse on heterogeneous clients
- Some clients may have been able to train **better local models**
  - Thus receiving **no benefit** from participating in FL
- Fair FL**
  - Focuses on clients with high losses during FL aggregation
  - Reduces the variance of the FL model accuracy over clients
- Personalisation:**
  - Fine-tune the FL model locally
  - Use **KD**, **EWC**, or FreezeBase (**FB**) as regularisers



**Personalisation-aware FL (PaFL)**

- Fair FL tends to flatten the top of the performance distribution
  - Has difficulties handling clients with very good local models
- PaFL uses regularisers like KD, EWC, FB **during** FL training
  - Allows their weight to **vary across rounds**
  - Maintains the **average** and **peak** performance of the model
    - While allowing training on heterogeneous clients
  - Can pre-empt the personalisation loss used after training
  - Inspired by Quantisation-aware training

**Experimental Results**

- Fair FL** is capable of reducing the variance of the accuracy distribution over clients
  - This leads to decreased performance on certain clients which are capable of training a very high quality local model
  - Thus, it **hurts** the **relative accuracy** distribution by skewing it towards negative values
  - It shows **no benefits** to later personalisation
- Personalisation-aware Federate Learning (PaFL)** provides an alternative which:
  - Allows training on heterogeneous clients while maintaining performance on the federated distribution
  - Leads to equivalent or **higher** accuracy on a centralised test-set representative of the federated distribution
  - May **reduce** the number of underperforming clients by up to **50%** and thus the need for personalisation