Models of fairness in federated learning

Kate Donahue and Jon Kleinberg

Department of Computer Science, Cornell   Departments of Information Science and Computer Science, Cornell

Federated learning

Consider an agent trying to solve a learning problem. It could use local learning (using only its local data), or *federated learning*, where multiple agents each learn models separately and combine parameters to form a global model.

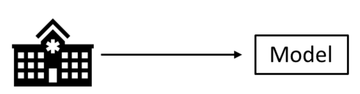

Figure 1: Local learning: Build a model based on its n_j data points.

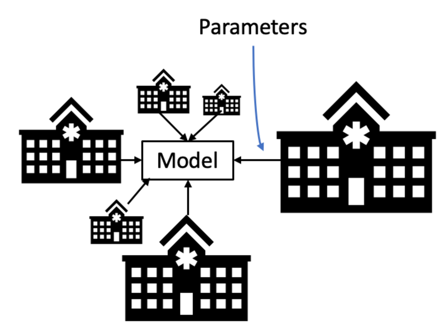

Figure 2: Federation: Build a model with $M-1$ other agents, each with n_i for $i \in [M]$ data points. Combine model parameters learned on local data.

Federated learning presents a bias-variance trade-off: With federation, there's more data (lower variance), but agents could differ from each other (higher bias).
Hospitals have different *sizes*: Larger hospitals skew the combined model: lower bias, so lower error. Smaller hospitals see higher bias, so higher error.

Our prior work: Model-sharing Games, Stability, and Optimality

Previously, we have analyzed federated learning in two related papers. In this work, we use the model first proposed in "Model-sharing Games", but ask questions and derive results completely independent of these prior works.

Model-sharing Games: Analyzing Federated Learning Under Voluntary Participation Donahue and Kleinberg, AAAI '21, `https://arxiv.org/abs/2010.00753`
In this first paper on federated learning, we use a lens of cooperative game theory to analyze which federating structures will be *stable* - where no agent wishes to move from its current federating coalition to another coalition (that would accept it). We derive exact expected MSE values for two federating games (linear regression and mean estimation) and consider three models of federation, offering varying levels of customization (vanilla federation, coarse-grained, and fine-grained). For each setting, we derive constructive examples of stable partitions.

Optimality and Stability in Federated Learning: A Game-theoretic Approach
Donahue and Kleinberg, Neurips '21, `https://arxiv.org/abs/2106.09580`
In this paper, we consider the relationship between a federating structure's stability and its optimality (how low the overall average error is). After giving an efficient, constructive algorithm for calculating an optimal arrangement, we show that optimal arrangements are not always stable, but that the worst stable solution has a cost no more than 9 times that of an optimal arrangement.

This work: federated learning and fairness

In this work, we consider two notions of *fairness* in federated learning. Both of these have been described in multiple other papers in the federated learning literature, but to our knowledge, we are the first paper to formally compare them.
Egalitarian fairness:
The motivating idea is that agents should have error rates that are roughly the same (*error ratio* not too high). One motivation for using this notion of fairness is that agents may have fewer samples for unfair historical reasons: we want to correct for this.

$$\frac{err_i(C)}{err_j(C)} \leq \lambda \ \forall i, j \in C$$

In this work, for egalitarian fairness, we prove a tight bound on error ratio for agents federating together.
Proportional (collaborative) fairness:
The motivating idea is that agents should have error rates that are (inversely) proportional to samples contributed. One motivation for using this notion of fairness is that we may wish to reward agents who contribute more data.

$$\frac{err_i(C)}{err_j(C)} \approx \frac{n_j}{n_i} \ \forall i \in C$$

In this work, for proportional fairness, we give conditions where a variant of proportional fairness is satisfied.

Egalitarian fairness

The main result for egalitarian fairness is Theorem 2, which bounds error ratio between any two players that are federating together, as a function of the size of the larger player.

Theorem 2. *Any "modular" federation method satisfies the error ratio bound:*

$$\frac{err_i(C)}{err_j(C)} \leq 2 \cdot c + 1 \ \forall i, j \in C$$

given largest player size $\leq c \cdot r$ *(for* r *noise/bias ratio).*

Theorem 3. *This bound is tight (up to an additive factor of* ϵ*).*

Definition 1 defines what we mean by "modular". In the full paper, we show that at least two federation methods (vanilla federation and fine-grained federation) are modular.

Definition 1. *Consider any coalition* C *with players* $s, \ell \in C$ *such that* $n_s \leq n_\ell$*. Then, a federating method is called modular if it satisfies the following four properties:*
Property 1: *The worst-case situation for the error ratio (the ratio of the small player's error to the large player's error) is always in the two-player case.*
Property 2: *The error ratio increases as the large player gets more samples.*
Property 3: *The error ratio decreases as the small player gets more samples.*
Property 4: *As* $\frac{n_s}{n_\ell} \rightarrow 0$*, the error ratio converges to the following fraction:* $\frac{\frac{r}{n_\ell}+2\sigma^2}{\frac{r}{n_\ell}}$.

Proportional fairness

Theorem 4. *Optimal fine-grained federation always displays sub-proportionality of error.*

For proportional fairness, our next result shows that vanilla federation could, in general, result in either sub- or super-proportional fairness.

Lemma 1. *There exist cases where a coalition using vanilla federation satisfies proportional, strict sub-proportional, and strict sup-proportional error.*

Table 1 gives an example where two players federating together results in sub-proportional error for the small player, while Table 2 gives an example where two players federating together results in super-proportional error for the small player. However, Theorem 5, below, shows that any federating arrangement with super-proportional error must be inherently unstable. Specifically, it must fail to be individually rational: at least one player would prefer to move to local learning.

Theorem 5. *Any individually rational coalition using vanilla federation satisfies sub-proportionality of errors.*

Examples

Coalition C	$err_s(C)$	$err_l(C)$	$\frac{err_s(\{n_s,n_\ell\})}{err_l(\{n_s,n_\ell\})}$	$2\cdot c+1$	$\frac{n_\ell}{n_s}$
$\{\{n_s\},\{n_\ell\}\}$	1.67	0.5	3.33	5	3.33
$\{n_s, n_\ell\}$	1.57	0.49	3.20	5	3.33

Table 1: Example of two-player game with $n_s = 6, n_\ell = 20$, with $r = 10$. Here, federating in the grand coalition $\{n_s, n_\ell\}$ is individually stable (and thus optimal, for weighted error). The grand coalition satisfies egalitarian fairness ($2 \cdot c + 1$ bound) and proportional fairness.

Coalition C	$err_s(C)$	$err_l(C)$	$\frac{err_s(\{n_s,n_\ell\})}{err_l(\{n_s,n_\ell\})}$	$2\cdot c+1$	$\frac{n_\ell}{n_s}$
$\{\{n_s\},\{n_\ell\}\}$	1.67	0.250	6.67	9	6.67
$\{n_s, n_\ell\}$	1.73	0.251	6.89	9	6.67

Table 2: Example of two-player game with $n_s = 6, n_\ell = 40$, with $r = 10$. Here, the grand coalition fails to be stable, as local learning minimizes weighted error. The grand coalition satisfies egalitarian fairness ($2 \cdot c + 1$ bound) but does not satisfy proportional fairness.

Future work

Future work could explore two main avenues:

- Other notions of fairness: for example, an additive version of egalitarian fairness, or proportional fairness based on the quality of data points, rather than quantity.
- Other federation methods, beyond vanilla and fine-grained.

Github: `https://github.com/kpdonahue/model_sharing_games`