Fairness and Incentives for Data Sharing and Collaborative Learning

Which notion of "fairness"? Why? How?

Xinyi Xu
https://xinyi-xu.com

Talk at Federated Learning One World (FLOW) Seminar
March 6, 2024

Agenda

- Motivation and high-level description for collaboration
- The fairness desideratum
- Gradient-Driven Rewards to Guarantee Fairness in Collaborative Machine Learning (NeurIPS’21)
- Fair yet Asymptotically Equal Collaborative Learning (ICML 23)
- Connection to our other works

Agenda

- Motivation and high-level description for collaboration
- The fairness desideratum
- Gradient-Driven Rewards to Guarantee Fairness in Collaborative Machine Learning (NeurIPS’21)
- Fair yet Asymptotically Equal Collaborative Learning (ICML 23)
- Connection to our other works

Why Collaborate?

- To share and leverage the **resources** (e.g., data, computational, infrastructural) of different parties (i.e., entities, research teams, organizations, companies, etc.) in a **mutually beneficial** way.
- Motivating use-cases (and interest/attention)
  - In different fields: medicine,[1] finance,[2] agriculture,[3] environment,[4] technology.[5]
  - In different sectors: industry,[6] academia, governmental agencies,[7], public sectors.[8]

1. Sheller, M.J., Edwards, B., Reina, G.A. et al. Federated learning in medicine: facilitating multi-institutional collaborations without sharing patient data. Sci Rep 10, 12598 (2020).
2. Angel N. B., et al. from IBM. Building privacy-preserving federated learning to help fight financial crime. 2023.
3. Krista Rizman Žalik and Mitja Žalik. A Review of Federated Learning in Agriculture. PubMed. 2023.
4. Y. Gao, L. Liu, B. Hu, T. Lei and H. Ma, "Federated Region-Learning for Environment Sensing in Edge Computing System," in *IEEE Transactions on Network Science and Engineering.* 2020.
5. McMahan B., et al., Communication-Efficient Learning of Deep Networks from Decentralized Data. In AISTATS 2017.
6. Paulik M., et al., at Apple. Federated Evaluation and Tuning for On-Device Personalization: System Design & Applications. 2022.
7. PETs Prize Challenge: Advancing Privacy-Preserving Federated Learning. NASA.
8. Data Management and Sharing Policy. NIH. 2023

Data sharing & collaborative learning

With multiple parties interacting with each other,

How to incentivize the parties to share and collaborate?

Incentives

What if these parties are self-interested and possibly competitive?

Fairness

Agenda

- Motivation and high-level description for collaboration
- The fairness desideratum
- Gradient-Driven Rewards to Guarantee Fairness in Collaborative Machine Learning (NeurIPS’21)
- Fair yet Asymptotically Equal Collaborative Learning (ICML 23)
- Connection to our other works

Fairness

But which notion of fairness?

Shapley value (SV)[1]

The more you contribute, the higher your SV is.

Egalitarian[2]

Everyone should be treated equally, or as much as possible.

Pigou-Dalton Principle (not covered in this talk)[3]

Given two outcomes of multiple parties, that differ in only one aspect, the more equitable outcome is preferred.

1. Lloyd Shapley. A value for n-person games. In H. W. Kuhn and A. W. Tucker, editors, Contributions to the Theory of Games, volume 2, pages 307–317. Princeton Univ. Press, 1953.
2. “An egalitarian favors equality of some sort: People should get the same, or be treated the same, or be treated as equals, in some respect.” https://plato.stanford.edu/entries/egalitarianism/
3. “It says that, all other things being equal, a social welfare function should prefer allocations that are more equitable.” https://en.wikipedia.org/wiki/Pigou%E2%80%93Dalton_principle

How to guarantee fairness?

With fair valuation and well-designed incentives.

First try: pay the parties with money.

Pros: conceptually very simple. Numerical/real values are very easy to handle (linear).

Cons: practically challenging. (i) who provides the funding; (ii) what is the denomination; (iii) how to actually make the transfer, etc.

Exploit the modes of learning and collaboration to identify a "**resource**" as incentive instead of money.

Setting for federated learning

For n self-interested parties, each with a local dataset $\mathscr{D}_i$.

The federated objective is

$$\omega^* := \text{argmin}_\omega \sum_i p_i F(\omega; \mathscr{D}_i).$$

In iteration t:

For party i

$$\Delta\omega_{i,t} \leftarrow -\eta \nabla_\omega F(\omega_{i,t}; \mathscr{D}_i) \; .$$

$$\omega_{i,t+1} \leftarrow \omega_{i,t} + u_{N,t} \; .$$

For server

$$u_{N,t} \leftarrow \sum_i p_i \Gamma \frac{\Delta\omega_{i,t}}{\|\Delta\omega_{i,t}\|} \; .$$

p_i is an importance coefficient, $F(\omega; \mathscr{D})$ is the loss, Γ is a normalizing constant.

Agenda

- Motivation and high-level description for collaboration
- The fairness desideratum
- Gradient-Driven Rewards to Guarantee Fairness in Collaborative Machine Learning (NeurIPS'21)
- Fair yet Asymptotically Equal Collaborative Learning (ICML 23)
- Connection to our other works

Gradient-Driven Rewards to Guarantee Fairness in Collaborative Machine Learning

Xinyi Xu, Lingjuan Lyu, Xingjun Ma, Chenglin Miao, Chuan-Sheng Foo, Bryan Kian Hsiang Low

In NeurIPS 2021

Gradients as the **resource**

In iteration t:

For party i

$$\boxed{\Delta\omega_{i,t}} \leftarrow -\eta \nabla_\omega F(\omega_{i,t}; \mathscr{D}_i) \,.$$

$$\omega_{i,t+1} \leftarrow \omega_{i,t} + \boxed{u_{N,t}} \,.$$

For server

$$\boxed{u_{N,t}} \leftarrow \sum_i p_i \Gamma \frac{\Delta\omega_{i,t}}{\|\Delta\omega_{i,t}\|} \,.$$

Are all the uploaded gradients $\Delta\omega_{i,t}$ of the same quality/**equally valuable**?

Is it **fair** that all parties receive the same downloaded gradient $u_{N,t}$?

p_i is an importance coefficient, $F(\omega; \mathscr{D})$ is the loss, Γ is a normalizing constant.

Fairly <u>valuing</u> gradients

Def. 1 Cosine Gradient Shapley value (CGSV).

$$\phi_i := \frac{1}{n} \sum_{S \subseteq [n] \setminus \{i\}} \frac{|S|!(n - |S| - 1)!}{n!} (\nu(S \cup \{i\}) - \nu(S))$$

With $\nu := \cos(u_S, u_N)$ where $u_i := \Gamma \frac{\Delta\omega_i}{\|\Delta\omega_i\|}$, $u_S := \sum_{i \in S} p_i u_i$.

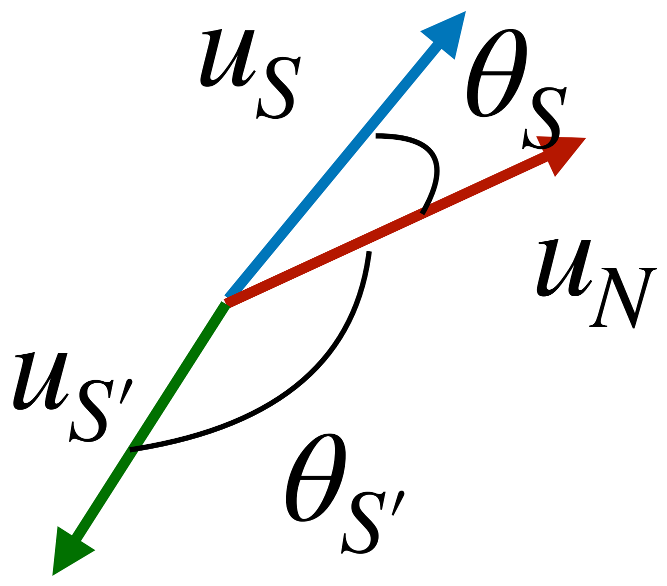

$$\nu(S) = \cos(u_S, u_N) = \cos(\theta_S)$$

$$\nu(S') = \cos(u_{S'}, u_N) = \cos(\theta_{S'})$$

Intuition for cosine similarity: If a gradient $\Delta\omega_{i,t}$ points to the right direction (of u_N), then it is valuable.

<u>Fairly</u> valuing gradients

Def. 1 Cosine Gradient Shapley value (CGSV).

$$\phi_i := \frac{1}{n} \sum_{S \subseteq [n] \setminus \{i\}} \frac{|S|!(n - |S| - 1)!}{n!} (\nu(S \cup \{i\}) - \nu(S))$$

The Shapley value has nice properties that formalize fairness, such as

<u>Null player</u>: if a party uploads non-valuable gradients, then the CGSV is 0. Effective against free-riders.

<u>Symmetry</u>: if two parties upload equally valuable gradients, then their CGSVs are equal.

The CGSV ϕ_i provides a fair value for gradient $\Delta\omega_{i,t}$ in the high-dimensional model parameter space.

Fair incentives

For party i, instead of the original $u_{N,t}$,

give the following (as downloaded gradient)

$$v_{i,t} \leftarrow \text{mask}(u_{N,t}, q_{i,t})$$

where $q_{i,t}$ is a cumulative function of $\phi_{i,t}$ and is monotonically increasing.

Intuition: A **higher** CGSV (contribution) gets a **better** downloaded gradient.

For party i, instead of

$$\omega_{i,t+1} \leftarrow \omega_{i,t} + u_{N,t} .$$

It performs:

$$\omega_{i,t+1} \leftarrow \omega_{i,t} + v_{i,t} .$$

Experiments

- *Fairness* metric
  - Pearson correlation coefficient between *standalone performance* & *final individual model performance*
  - The standalone performance is the surrogate for the parties' values/contributions where the final individual model performance (from FL) represent the parties' incentives.
  - A high Pearson correlation between them indicates good fairness.
- *Accuracy* metric
  - Maximum and average final test accuracies over all parties.

Experiments

- Baselines
  - FedAvg[1]
  - q-FFL[2], CFFL[3]
  - Shapley-value based[4]
  - Euclidean distance variant

- Data partitions
  - uniform (UNI)
  - powerlaw (POW)
    - Individual datasets of different **sizes**
  - classimbalance (CLA)
    - Individual datasets with different **available classes**

e.g. MNIST, for n=5, the parties have {1,3,5,7,10} classes respectively

1. Communication-Efficient Learning of Deep Networks from Decentralized Data. H. Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, Blaise Agüera y Arcas, 2017, AISTATS.
2. Fair Resource Allocation in Federated Learning. Tian Li, Maziar Sanjabi, Ahmad Beirami, Virginia Smith. 2020, ICLR.
3. Collaborative fairness in federated learning. Lingjuan Lyu, Xinyi Xu, Qian Wang. 2020, LNCS.
4. Profit Allocation for Federated Learning. Tianshu Song, Yongxin Tong, Shuyue Wei, IEEE Big Data, 2019.

Fairness results

	MNIST						CIFAR-10			MR	SST
No. Agents	10			20			10			5	5
Data Partition	UNI	POW	CLA	UNI	POW	CLA	UNI	POW	CLA	POW	POW
FedAvg	-45.60	55.24	24.12	0.85	-32.58	40.83	18.47	97.48	98.75	48.68	57.50
q-FFL	-44.73	39.00	22.38	-22.01	38.71	48.07	-17.64	51.33	94.06	56.43	-75.92
CFFL	83.57	91.80	81.24	82.52	94.70	85.71	78.25	72.55	81.31	96.85	93.34
ECI	85.26	**99.83**	**99.98**	80.95	**99.41**	95.21	75.85	79.50	99.55	97.69	95.00
DW	89.15	98.93	65.34	86.94	99.63	35.21	-23.14	91.97	45.45	**99.20**	97.12
RR	83.77	71.17	-26.75	-18.64	25.47	95.86	30.67	0.70	90.67	44.16	-25.11
Ours (EU)	84.25	98.25	99.82	80.55	97.77	**99.97**	78.25	94.24	94.95	97.58	93.21
Ours ($\beta = 1$)	94.03	95.74	94.54	84.47	96.39	97.23	**98.80**	**98.78**	**99.89**	96.01	98.20
Ours ($\beta = 1.2$)	94.75	97.28	96.23	90.52	97.72	95.21	91.07	91.59	99.82	96.12	**98.47**
Ours ($\beta = 1.5$)	**96.34**	86.99	95.37	82.68	90.94	98.75	93.55	93.78	95.89	95.32	97.88
Ours ($\beta = 2$)	94.66	91.20	95.38	**96.90**	91.33	94.32	89.80	88.78	93.39	92.22	95.74

Accuracy results

	MNIST						CIFAR-10			MR	SST
No. Agents	10			20			10			5	5
Data Partition	UNI	POW	CLA	UNI	POW	CLA	UNI	POW	CLA	POW	POW
Standalone	91 (91)	88 (92)	53 (92)	91 (91)	89 (92)	48 (90)	46 (47)	43 (49)	31 (44)	47(56)	31(34)
FedAvg	93 (94)	92 (94)	53 (93)	93 (93)	92 (94)	49 (92)	48 (48)	47 (50)	32 (47)	51(63)	33(35)
q-FFL	85 (91)	27 (45)	44 (64)	88 (91)	48 (53)	40 (59)	41 (46)	36 (36)	22 (28)	12(18)	23(25)
CFFL	90 (92)	85 (90)	34 (44)	91 (93)	88 (91)	39 (46)	39 (41)	35 (45)	22 (40)	44(53)	31(32)
ECI	94 (94)	92 (94)	53 (94)	94 (94)	92 (94)	49 (92)	49 (49)	47 (51)	31 (46)	56(61)	33(34)
DW	93 (94)	92 (94)	53 (93)	93 (93)	92 (94)	49 (92)	48 (48)	47 (50)	32 (47)	51(62)	33(35)
RR	94 (95)	**95** (95)	64 (72)	94 (95)	94 (95)	50 (56)	47 (59)	49 (51)	26 (29)	**63**(65)	**36**(36)
Ours (EU)	94 (94)	94 (94)	54 (94)	94 (94)	94 (94)	49 (92)	49 (49)	49 (51)	32 (46)	54(59)	34(36)
Ours ($\beta = 1$)	96 (97)	94 (95)	74 (**95**)	95 (96)	96 (97)	65 (93)	61 (**62**)	**60** (**62**)	35 (**54**)	62(**76**)	35(36)
Ours ($\beta = 1.2$)	94 (95)	**95** (95)	**75** (**95**)	96 (96)	**96** (**97**)	65 (93)	61 (**62**)	**60** (**62**)	35 (**54**)	62(75)	34(**37**)
Ours ($\beta = 1.5$)	**97** (**97**)	**95** (95)	**75** (**95**)	**96** (**97**)	94 (95)	65 (93)	61 (**62**)	59 (**62**)	35 (**54**)	62(74)	35(**37**)
Ours ($\beta = 2$)	96 (96)	**95** (**96**)	73 (94)	**97** (**97**)	95 (96)	**66** (**95**)	**62** (**62**)	**61** (**62**)	**36** (**54**)	62(75)	35(**37**)

Average (maximum) test accuracies over all agents.

Discussion

- Fairness in incentives in action
    - Each party’s interest is protected, they are **fairly** incentivized based on their contributions measured in *(Cosine Gradient) Shapley values*;
    - High fairness without compromising accuracy.

Agenda

- Motivation and high-level description for collaboration
- The fairness desideratum
- Gradient-Driven Rewards to Guarantee Fairness in Collaborative Machine Learning (NeurIPS'21)
- Fair yet Asymptotically Equal Collaborative Learning (ICML 23)
- Connection to our other works

Fair yet Asymptotically Equal Collaborative Learning

Xiaoqiang Lin*, Xinyi Xu*, See-Kiong Ng, Chuan-Sheng Foo, Bryan Kian Hsiang Low

In ICML 2023

Fairness vs. Fairness

Can different notions of fairness "co-exist"?

The Shapley value: the incentives should be commensurate to the contribution or value.

Fairness

Vs.

The egalitarian: all parties should be treated equally or as much as possible.[1]

Equality

1. Tian Li, Maziar Sanjabi, Ahmad Beirami, Virginia Smith. Fair Resource Allocation in Federated Learning. ICLR 2020.

Why equality?

Because fairness may widen inequality.

- “Rich get richer”: widened “gaps” between nodes.

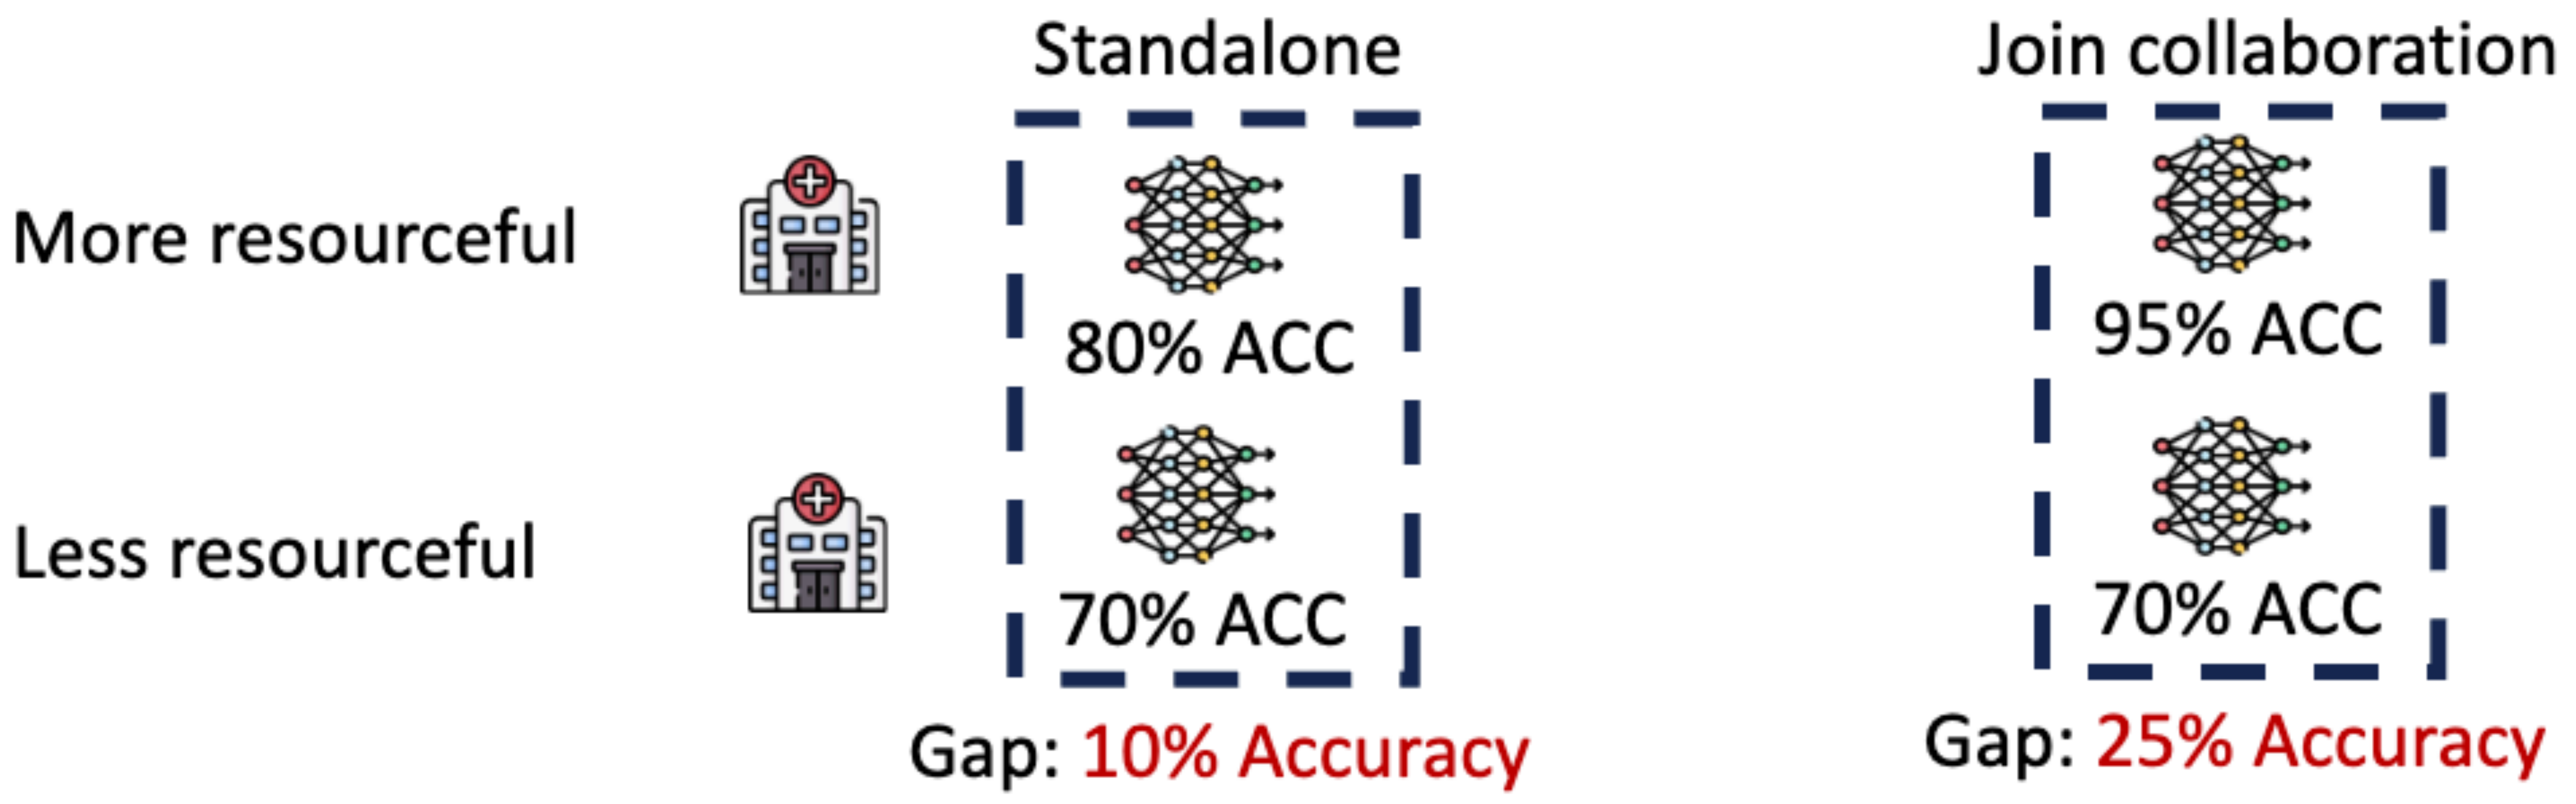

Fair yet asymptotically equal

The **fairness** is via the convergence rates at which the parties' models converge to the optimum:

"More valuable parties converge more quickly."

While the **asymptotic equality** is by guaranteeing that all parties are converging to the same (thus equal) optimum asymptotically:

"All parties converge to the same (model) eventually."

Fair yet asymptotically equal

Design and realize **fair in-training incentives** for the parties.

1. Collect the Shapley values (for each party), cumulatively; stop when they are no longer fluctuating.

2. Design incentives based on the Shapley values and ensure that all parties are equal **in the end (asymptotically)**.

Obtain Shapley values

Cumulative over iterations: $\psi_{i,t} := \frac{1}{t} \sum_{l=0}^{t} \phi_{i,t}$

Iteration t: $\boxed{\phi_{i,t} := \frac{1}{n} \sum_{S \subseteq [n] \setminus \{i\}} \binom{n-1}{|S|}^{-1} U(S \cup \{i\}) - U(S)}$

the Shapley value

Gradient dot product: $U(S) = \langle \Delta\omega_{S,t}, \Delta\omega_{N,t} \rangle$

Stopping criterion

Continue to calculate $\psi_{i,t} := \frac{1}{t}\sum_{l=0}^{t} \phi_{i,t}$ over iterations t,

$\boxed{\text{until there is } \textit{no evidence that it is fluctuating} \text{ (Proposition 1).}}$

Intuition: the more gradients (over iterations) we observe from the parties, the more accurately $\psi_{i,t}$ reflects their values, and once $\psi_{i,t}$ is sufficiently accurate, stop.

Incentive realization

The parties are "synchronized" with the global model with frequencies commensurate with their Shapley values.

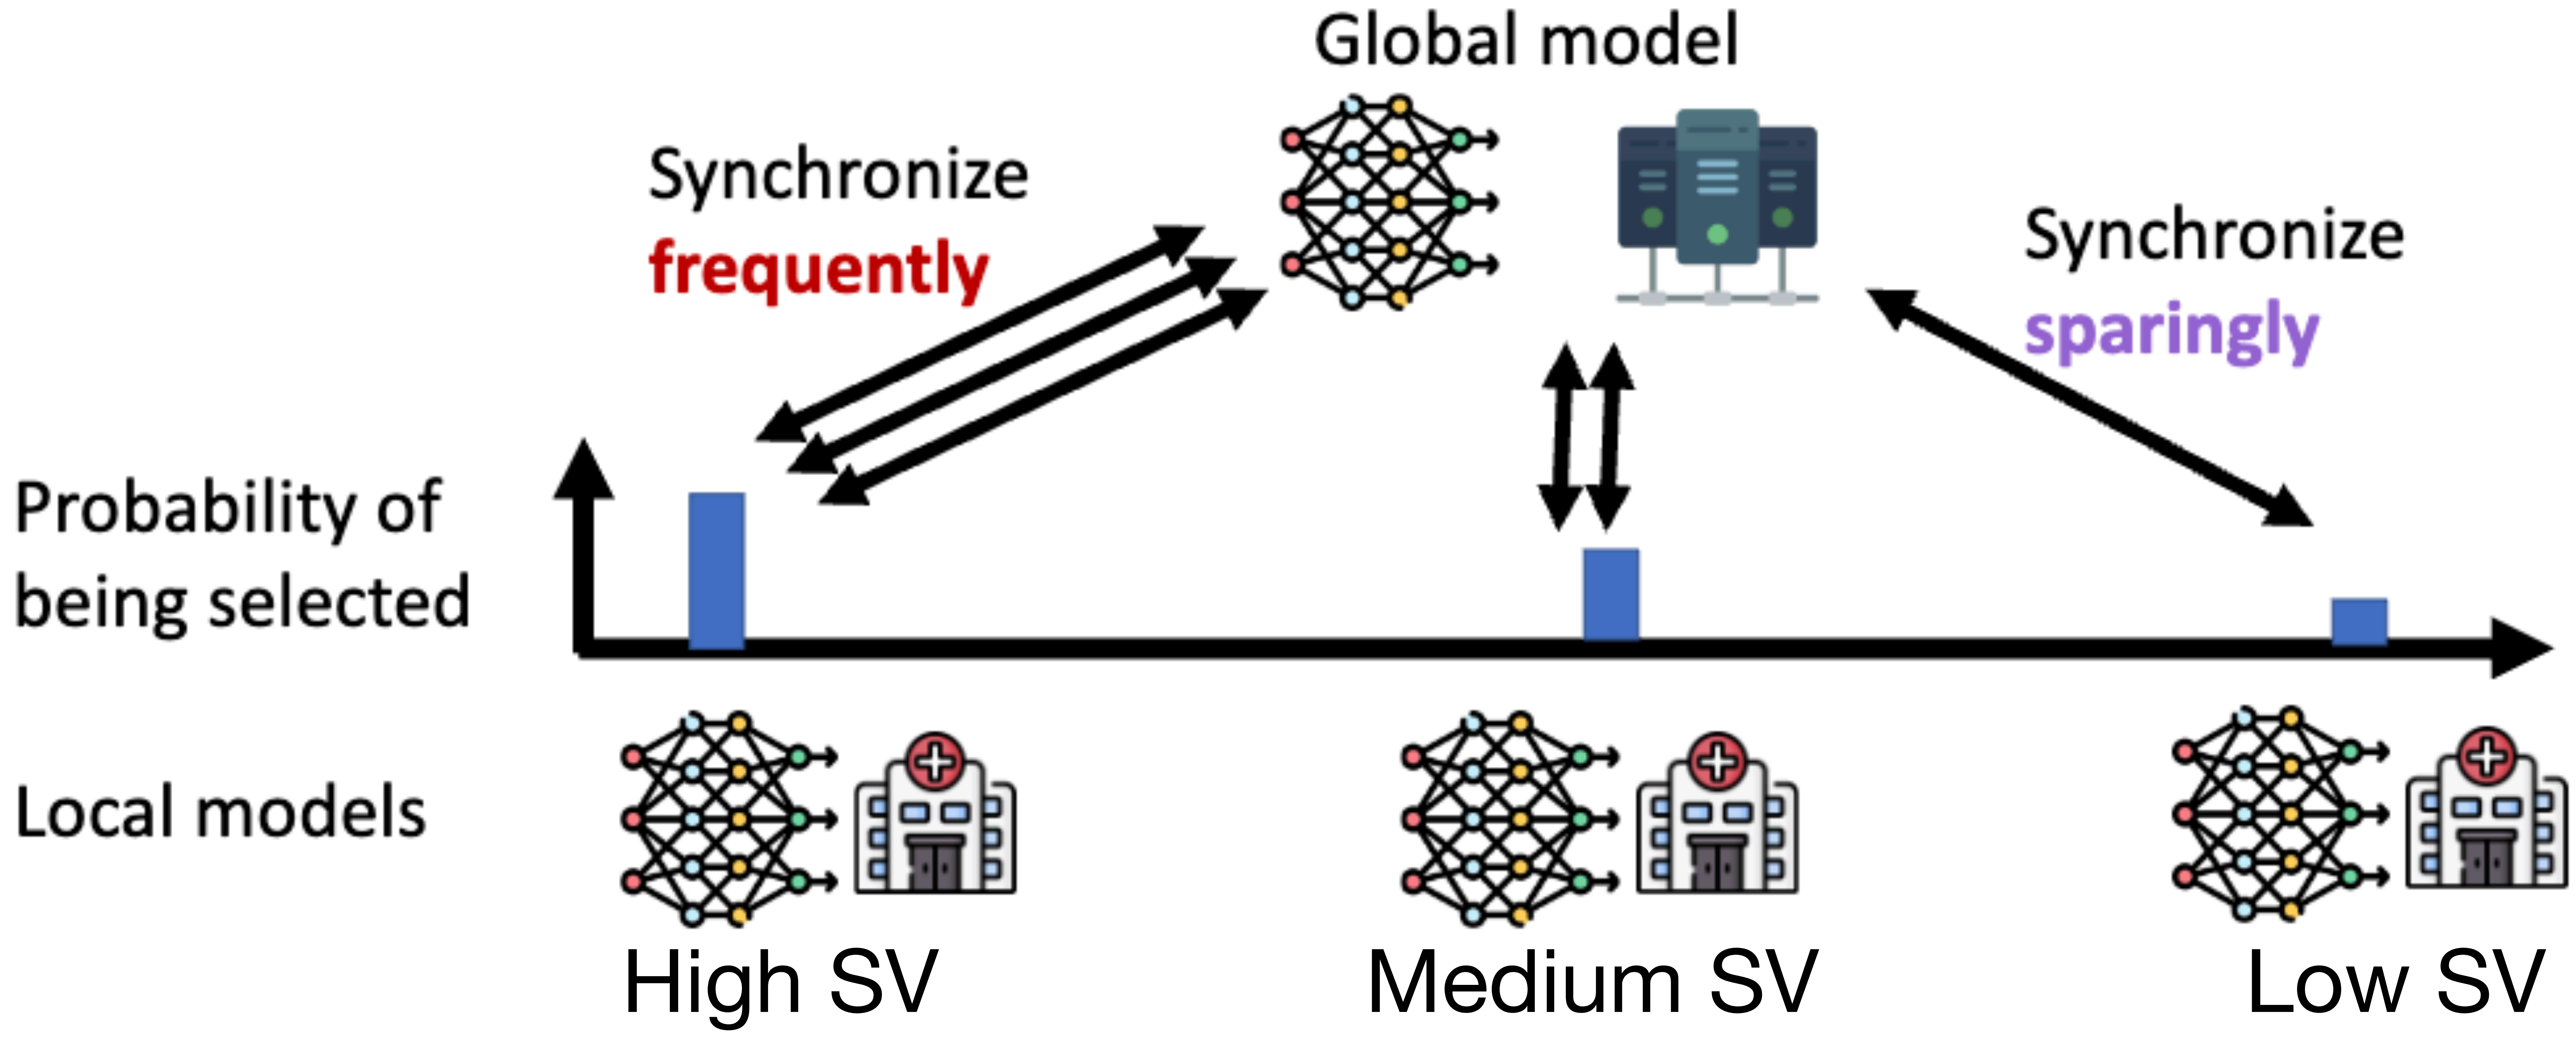

Incentive realization

The parties are “synchronized” with the global model with frequencies commensurate with their Shapley values.

Probability of party i being sampled in an iteration to synchronize its model:

$$\rho_i := \frac{\exp(\psi_{i,T_{\text{stop}}})/\beta}{\sum_{i' \in [n]} \exp(\psi_{i',T_{\text{stop}}})/\beta}$$

Fairness (Proposition 2): parties with higher SV are sampled more frequently, thus converge faster.

Equality (Proposition 3): *all* parties are sampled with non-zero probabilities, and converge to the equal optimum eventually.

Experiments - metrics

- Fairness metric:
  - Pearson coefficient between SVs and test performance.
- Equality metric:
  - Standard deviation of test performance over parties.
- Accuracy metric:
  - Maximum and average online test performance over parties.

Experiments - settings

- Federated online incremental learning
  - Training data are incrementally collected during training
- Federated reinforcement learning
  - Exploration trajectories (i.e., training data) of the environment incrementally collected during training
- Baselines:
  - FedAvg,[1] q-FFL,[2] FGFL,[3] GoG,[4] Standalone.

1. Communication-Efficient Learning of Deep Networks from Decentralized Data. H. Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, Blaise Agüera y Arcas, 2017, AISTATS.
2. Fair Resource Allocation in Federated Learning. Tian Li, Maziar Sanjabi, Ahmad Beirami, Virginia Smith. 2020, ICLR.
3. Gradient-driven rewards to guarantee fairness in collaborative machine learning. Xu et al. 2021. NeurIPS.
4. Game of gradients: Mitigating irrelevant clients in federated learning. Nagalapatti, L. and Narayanam, 2021. AAAI.

Experiments - performance

(a) Average of online accuracy

	MNIST	CIFAR-10	HFT	ELECTRICITY	PATH
FedAvg	0.483(0.019)	0.166(0.011)	0.499(0.046)	1.408(0.081)	0.255(0.004)
qFFL	0.101(0.011)	0.100(0.004)	0.281(0.079)	1.413(0.055)	0.101(0.011)
FGFL	0.485(0.018)	0.169(0.010)	0.496(0.056)	1.571(0.090)	0.154(0.009)
GoG	0.572(0.015)	0.193(0.006)	0.556(0.016)	1.394(0.032)	0.288(0.004)
Standalone	0.481(0.013)	0.153(0.004)	0.540(0.014)	1.581(0.083)	0.202(0.003)
Ours	**0.611(0.009)**	**0.195(0.007)**	**0.581(0.014)**	**0.139(0.002)**	**0.302(0.005)**

(b) Minimum of online accuracy

	MNIST	CIFAR-10	HFT	ELECTRICITY	PATH
FedAvg	0.478(0.020)	0.165(0.011)	0.497(0.046)	1.405(0.081)	0.250(0.005)
qFFL	0.101(0.011)	0.100(0.004)	0.281(0.079)	1.413(0.055)	0.101(0.011)
FGFL	0.437(0.030)	0.167(0.010)	0.493(0.058)	1.497(0.071)	0.153(0.008)
GoG	0.553(0.014)	0.189(0.006)	0.548(0.017)	1.383(0.032)	0.271(0.004)
Standalone	0.279(0.024)	0.131(0.006)	0.515(0.017)	1.281(0.065)	0.131(0.006)
Ours	**0.603(0.010)**	**0.193(0.007)**	**0.581(0.014)**	**0.139(0.002)**	**0.298(0.005)**

For ELECTRICITY, regression error is measured.

Experiments - fairness vs. equality

β	MNIST	CIFAR-10	HFT	ELECTRICITY	PATH
1/350	0.642(9.52e-03)	0.490(1.86e-03)	0.448(6.04e-05)	0.581(1.63e-03)	0.516(3.18e-03)
1/150	0.647(2.2e-03)	0.400(6.5e-04)	0.378(5.8e-05)	**0.676**(2.95e-04)	**0.557**(1.41e-03)
1/100	**0.705**(1.13e-03)	**0.507**(3.80e-04)	0.415(1.05e-04)	0.572(1.63e-04)	0.312(1.15e-03)
1/50	0.641(5.66e-04)	0.297(2.67e-04)	**0.476**(**8.88e-17**)	0.286(8.17e-05)	0.282(8.32e-04)
1/20	0.466(4.91e-04)	0.131(3.05e-04)	0.217(1.84e-06)	0.166(**5.15e-05**)	0.127(7.91e-04)
1/10	0.120(4.97e-04)	0.034(2.91e-04)	0.090(1.24e-05)	0.068(6.01e-05)	0.018(7.49e-04)
1	0.171(**3.79e-04**)	0.063(**2.43e-04**)	-0.187(1.30e-05)	0.193(6.11e-05)	0.082(**6.41e-04**)
1000	-0.005(3.88e-04)	-0.185(2.77e-04)	-0.079(1.61e-05)	-0.034(5.18e-05)	0.157(7.19e-04)

Pearson correlation coefficient (standard deviation of performance over parties).

*A **high** Pearson correlation coefficient -> good fairness.*

*A **low** standard deviation -> good equality (or egalitarian fairness).*

Discussion

- Fair in-training incentives:
  - No waiting till the end for incentives -> parties can get “paid” sooner.
  - The incentives are realized within the FL framework -> no external resources (e.g., money) required.
  - Fairness without compromising accuracy.
  - Asymptotic equality without compromising fairness or accuracy.

Agenda

- Motivation and high-level description for collaboration
- The fairness desideratum
- Gradient-Driven Rewards to Guarantee Fairness in Collaborative Machine Learning (NeurIPS’21)
- Fair yet Asymptotically Equal Collaborative Learning (ICML 23)
- Connection to our other works

Our other works

	Desiderata or Fairness	Statistic	Mode of learning	Mode of collaboration
[1] NeurIPS 21	SV	Model updates/ gradients	Supervised	Federated learning
[2] ICML 23	**SV + Egalitarian**	Model updates/ gradients	Supervised	Federated learning
[3] ICML 22	SV	(Fisher information of) observations	Parametric estimation	**Joint Bayesian inference**
[4] AAAI 22	SV	(Generated/synthetic) data	**Unsupervised**	Data sharing
[5] AISTATS 23	**PDP + Individual Rationality**	Mutual information	**Active learning**	Joint data selection
[6] ICML 23	SV	Estimate of average treatment effect	**Causal inference**	Data sharing
[7] NeurIPS 23	SV	**(Predictions of) trained ML model**	Supervised	Model fusion
[8] NeurIPS 23	SV + **Privacy**	Sufficient statistic	Parametric estimation	Joint Bayesian inference

Bold represents a new consideration or extension.

Our other works

- [9] A formal framework for approximate SV fairness framework, AAAI 2023.

- [10] Survey on Data Valuation IJCAI 2022.

- Recent book chapters in Federated Learning: Theory and Practice: fairness (chapter 8), data valuation (chapter 15), and incentives (chapter 16).

Thank you! Questions?

References

1. Gradient Driven Rewards to Guarantee Fairness in Collaborative Machine Learning. Xinyi Xu, Lingjuan Lyu, Xingjun Ma, Chenglin Miao, Chuan Sheng Foo, and Bryan Kian Hsiang Low. In NeurIPS 2021.
2. Fair yet Asymptotically Equal Collaborative Learning. Xiaoqiang Lin, Xinyi Xu, See-Kiong Ng, Chuan Sheng Foo, and Bryan Kian Hsiang Low. In ICML 2023.
3. On the Convergence of the Shapley Value in Parametric Bayesian Learning Games. Lucas Agussurja, Xinyi Xu, and Bryan Kian Hsiang Low. In ICML 2022.
4. Incentivizing Collaboration in Machine Learning via Synthetic Data Rewards. Sebastian Sheng Hong Tay, Xinyi Xu, Chuan Sheng Foo, and Bryan Kian Hsiang Low. In AAAI 2022.
5. FAIR: Fair Collaborative Active Learning with Individual Rationality for Scientific Discovery. Xinyi Xu, Zhaoxuan Wu, Arun Verma, Chuan Sheng Foo, and Bryan Kian Hsiang Low. In AISTATS 2023.
6. Collaborative Causal Inference with Fair Incentives. Rui Qiao, Xinyi Xu, and Bryan Kian Hsiang Low. In ICML 2023.
7. Model Shapley: Equitable Model Valuation with Black-box Access. Xinyi Xu, Thanh Lam, Chuan-Sheng Foo, and Bryan Kian Hsiang Low. In NeurIPS 2023.
8. Incentives in Private Collaborative Machine Learning. Rachael Hwee Ling Sim, Yehong Zhang, Trong Nghia Hoang, Xinyi Xu, Bryan Kian Hsiang Low, and Patrick Jaillet. In NeurIPS 2023.
9. Probably Approximate Shapley Fairness with Applications in Machine Learning. Zijian Zhou, Xinyi Xu, Rachael Hwee Ling Sim, Chuan Sheng Foo, and Bryan Kian Hsiang Low. In AAAI 2023.
10. Data Valuation in Machine Learning: "Ingredients", Strategies, and Open Challenges. Rachael Hwee Ling Sim, Xinyi Xu, and Bryan Kian Hsiang Low. In IJCAI 2022, survey track.