The Fair Value of Data Under Heterogeneous Privacy Constraints in Federated Learning

BLISS Berkeley Laboratory for Information and System Sciences

Justin Kang [1] Ramtin Pedarsani [2] Kannan Ramchandran [1]

[1]UC Berkeley [2]UC Santa Barbara

Berkeley UNIVERSITY OF CALIFORNIA

Introduction

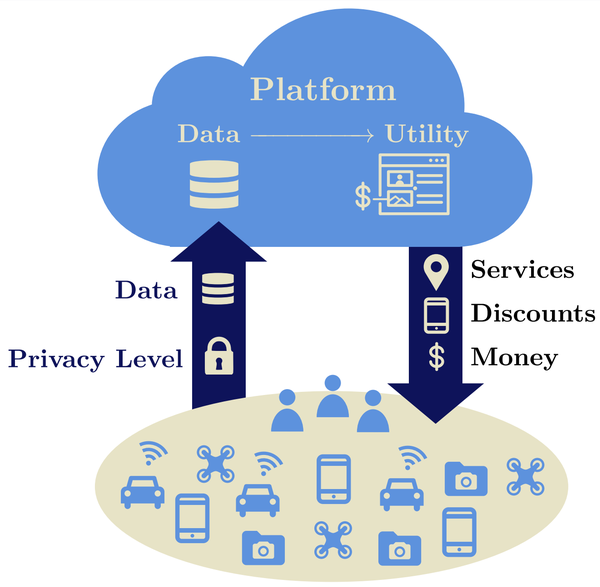

- Platforms collect data and do not fairly compensate users.
- Individuals are becoming more conscious of their privacy.
- In response, platforms are using privacy preserving tools:
  - Federated Learning
  - Differential Privacy

- In general private learning exhibits a trade-off: Privacy $\implies$ Less Utility.

What is a fair price to pay users for their data at a given privacy level?

Problem Setting

- Privacy level ϵ_i of a user can be defined via Differential Privacy:

$$\Pr(A(\mathbf{x}) \in S) \leq e^{\epsilon_i} \Pr(A(\mathbf{x}') \in S).$$

  Where $\epsilon_i = 0$ represents the case where no data is provided. Similarly Federated Learning can also be used (see example).
- Users have some privacy sensitivity $c_i(\epsilon_i)$. They seek to maximize the difference between their payment $\phi_i(\boldsymbol{\epsilon})$ and privacy sensitivity:

$$u_i(\boldsymbol{\epsilon}) = \phi_i(\boldsymbol{\epsilon}) - c_i(\boldsymbol{\epsilon}).$$

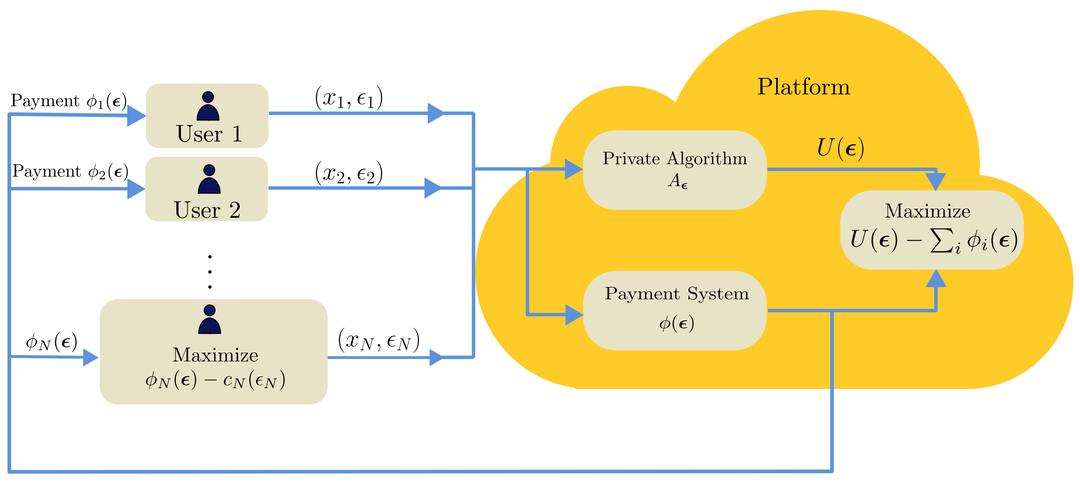

- Platform knows function $U(\boldsymbol{\epsilon})$, which describes it's privacy-utility trade-off. Seeks to maximize the difference $U(\boldsymbol{\epsilon}) - \sum \phi_i(\boldsymbol{\epsilon})$.

1. Given a payment system, how can we benchmark its fairness?
2. If we want to impose a regulatory constraint of fairness, how can we derive rules that a mechanism designer should satisfy?

Fairness in the Context of Privacy: Shapley Value

Inspired by powerful concept of Shapley Value from coalitional game theory, we consider some axioms that we want the payment of each agent to satisfy.

- **Axiom 1:** The sum of all the values $\phi_i(\boldsymbol{\epsilon})$ at privacy level $\boldsymbol{\epsilon}$ is $U(\boldsymbol{\epsilon})$.
- **Axiom 2:** If $U(\boldsymbol{\epsilon}) = V(\boldsymbol{\epsilon}) + W(\boldsymbol{\epsilon})$, then $\phi_i^U(\boldsymbol{\epsilon}) = \phi_i^W(\boldsymbol{\epsilon}) + \phi_i^V(\boldsymbol{\epsilon})$.
- **Axiom 3:** Two agents have the same marginal contribution when added to any coalition of agents $\implies$ $\phi_i(\boldsymbol{\epsilon}) = \phi_j(\boldsymbol{\epsilon})$.

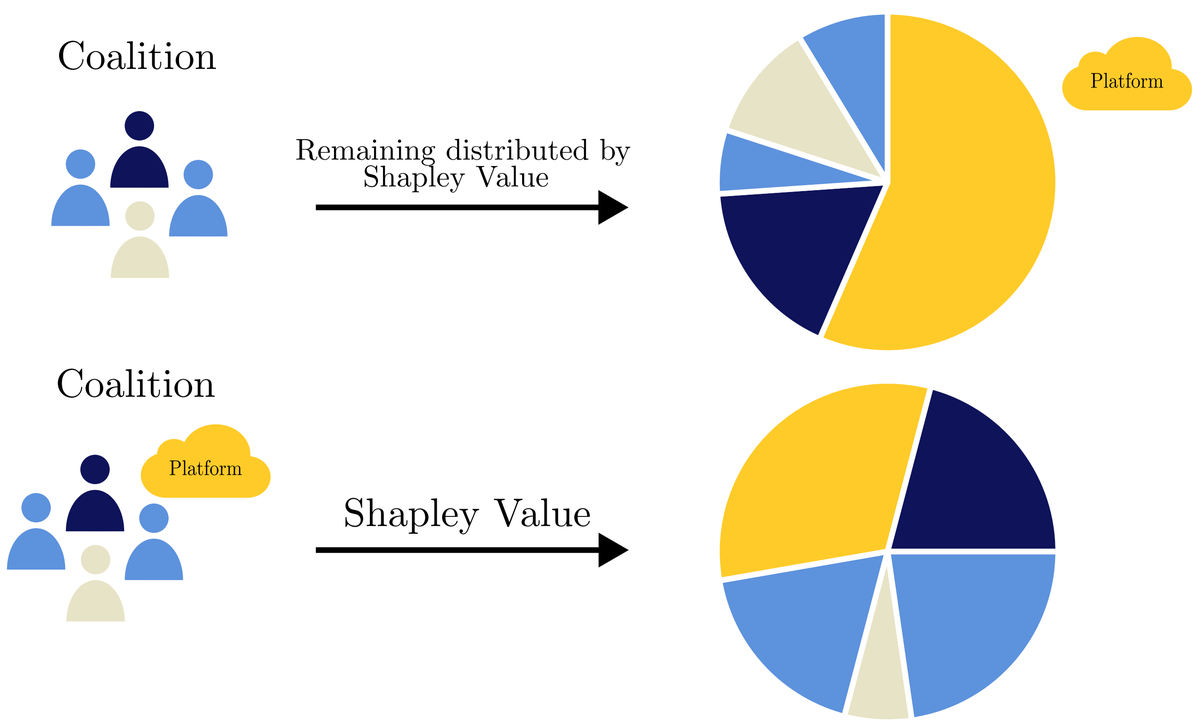

We consider two ways of defining fairness, both with their merits:

1. The platform keeps $(1 - \alpha(\boldsymbol{\epsilon}))U(\boldsymbol{\epsilon})$, and the remaining utility is distributed according to values that satisfy axioms 1-3.
2. Axioms 1-3 where the platform is viewed as a member of the coalition.

Key Takeaway: Notions of Fairness

Using Axioms 1-3 between the users the fair value is of the form:

$$\phi_p(\boldsymbol{\epsilon}) = (1 - \alpha(\boldsymbol{\epsilon}))U(\boldsymbol{\epsilon})$$

$$\phi_i(\boldsymbol{\epsilon}) = \alpha(\boldsymbol{\epsilon}) \cdot \underbrace{\frac{1}{N} \sum_{S \subseteq [N] \setminus \{i\}} \frac{1}{\binom{N-1}{|S|}} \Big(U(\boldsymbol{\epsilon}_{S \cup \{i\}}) - U(\boldsymbol{\epsilon}_S) \Big)}_{\text{weighted average marginal contribution}},$$

Platform can optimize over $\alpha(\boldsymbol{\epsilon})$, creating a mechanism design problem.

Axioms 1-3 between the users and the platform, unique fair value:

$$\phi_p(\boldsymbol{\epsilon}) = \frac{1}{N+1} \cdot \underbrace{\sum_{S \subseteq [N]} \frac{1}{\binom{N}{|S|}} U(\boldsymbol{\epsilon}_S)}_{\text{weighted sum over coalitions}},$$

$$\phi_i(\boldsymbol{\epsilon}) = \frac{1}{N+1} \cdot \underbrace{\sum_{S \subseteq [N] \setminus \{i\}} \frac{1}{\binom{N}{|S|+1}} \Big(U(\boldsymbol{\epsilon}_{S \cup \{i\}}) - U(\boldsymbol{\epsilon}_S) \Big)}_{\text{weighted sum of marginal contributions}}.$$

Serves as a principled benchmark for evaluating fairness.

Example: Fairness-Constrained Mechanism Design

$\boldsymbol{\epsilon}$-DP linear-Laplace mean estimator:

$$A(\mathbf{X}) = \mathbf{w}(\boldsymbol{\epsilon})^T \mathbf{X} + Z$$

where $N = 2$ samples X_1 and $X_2 \in \{-1, 1\}$ provided by each user, ϵ_i restricted to binary levels.

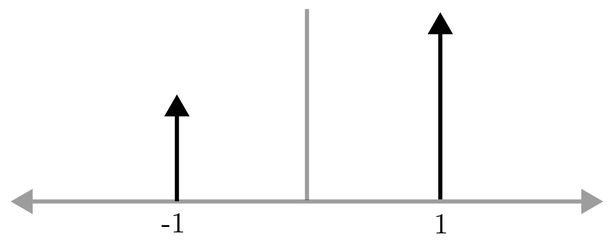

- Users experience a privacy cost c if they choose a lower privacy level:

$$u_i(\mathbf{p}_1, \mathbf{p}_2) = \mathbf{p}_1^T (\boldsymbol{\alpha} \odot \boldsymbol{\Phi}_i) \mathbf{p}_2 - [0 \; c]^T \mathbf{p}_i.$$

- Platform choose $\boldsymbol{\alpha}$ to maximize utility minus payments at equilibrium:

$$\max_{\boldsymbol{\alpha}} \quad \mathbf{p}_1^T \mathbf{U} \mathbf{p}_2 - \mathbf{p}_1^T (\boldsymbol{\alpha} \odot \mathbf{U}) \mathbf{p}_2$$
$$\text{s.t.} \quad (\mathbf{p}_1, \mathbf{p}_2) \in \text{NE}(\boldsymbol{\alpha}).$$

Sensitivity	Optimal Strategy
Low	Pay users enough so they choose minimal privacy
Medium	Pay users some amount so they play a mixed strategy
High	Pay users nothing

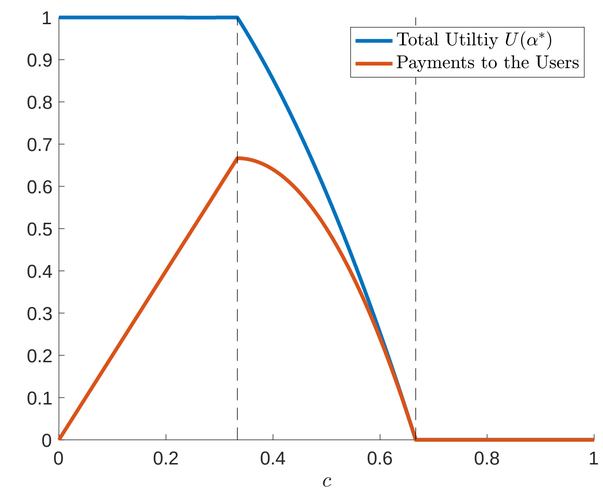

- Three optimal regions also proved to hold when $N > 2$ if $\boldsymbol{\alpha}$ constant.

Federated Mean Estimation

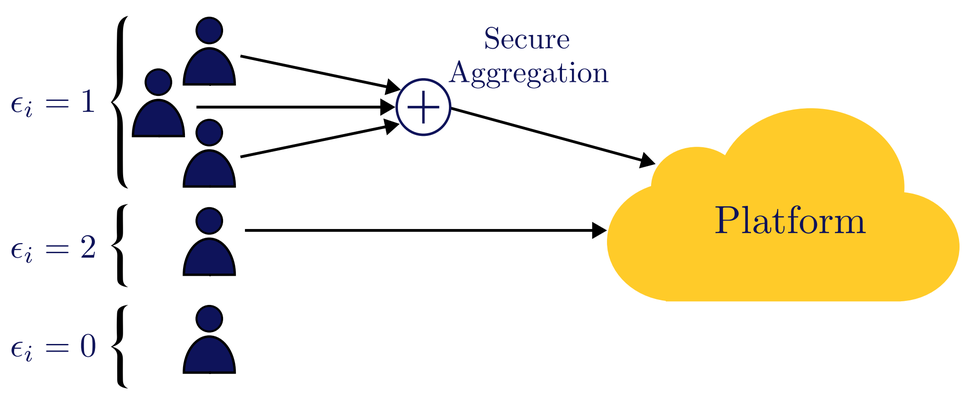

Federated Learning is compatible with the privacy level framework:

- $\epsilon_i = 2$: Users directly send data to the platform.
- $\epsilon_i = 1$: Users federate, securely aggregating their data (gradients, local models, etc.) before sending it to the platform.
- $\epsilon_i = 0$: Users do not to send their data to the platform.

Given the privacy-utility trade-off $U(\boldsymbol{\epsilon})$, the fair payments for either setting can be computed.

Further Reading

[1] Kang, Justin, Ramtin Pedarsani, and Kannan Ramchandran. "The Fair Value of Data Under Heterogeneous Privacy Constraints." Accepted, TMLR (2023)