Paid with Models:
Optimal Contract Design for Collaborative Machine Learning

Bingchen Wang*, Zhaoxuan Wu, Fusheng Liu, Bryan Kian Hsiang Low

*Corresponding author (bingchen@nus.edu.sg)

NUS National University of Singapore

GLOW.AI

The Herculean Task of Model Training

- **Motivation**
- Models as Rewards
- Optimal Contract
- Results

Geoffrey Hinton in conversation with Fei-Fei Li — Responsible AI development, Oct 4, 2023.

"Not a single university in the US today can train a ChatGPT in terms of the compute power.

And I think combining all universities of the US, A100 or H100—probably nobody has it, but A100 cannot train a ChatGPT."

Can we design mechanisms to democratize AI?

Collaborative Machine Learning (CML)

- **Motivation**
- Models as Rewards
- Optimal Contract
- Results

A promising crowdsourcing paradigm to democratize AI?

The incentive problem of CML

- **Motivation**
- Models as Rewards
- Optimal Contract
- Results

Conflict of interests

Resource contribution increases model performance but costs money.

Collective goal:
Maximize the model performance

Private goal:
Maximize the net profit

The incentive problem of CML

- **Motivation**
- Models as Rewards
- Optimal Contract
- Results

Conflict of interests

Karimireddy, Guo and Jordan (2022)

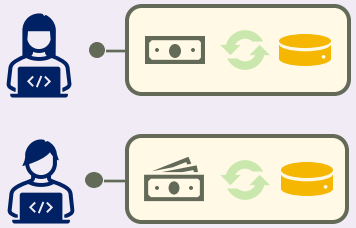

Parties have different contribution costs.

Catastrophic freeriding

The incentive problem of CML

Conflict of interests + **Complete Information**

- **Motivation**
- Models as Rewards
- Optimal Contract
- Results

Proposition 2 (Contract Design with Complete Information)

The optimal strategy is to offer participants the same best model but require them to contribute different amounts of resources.

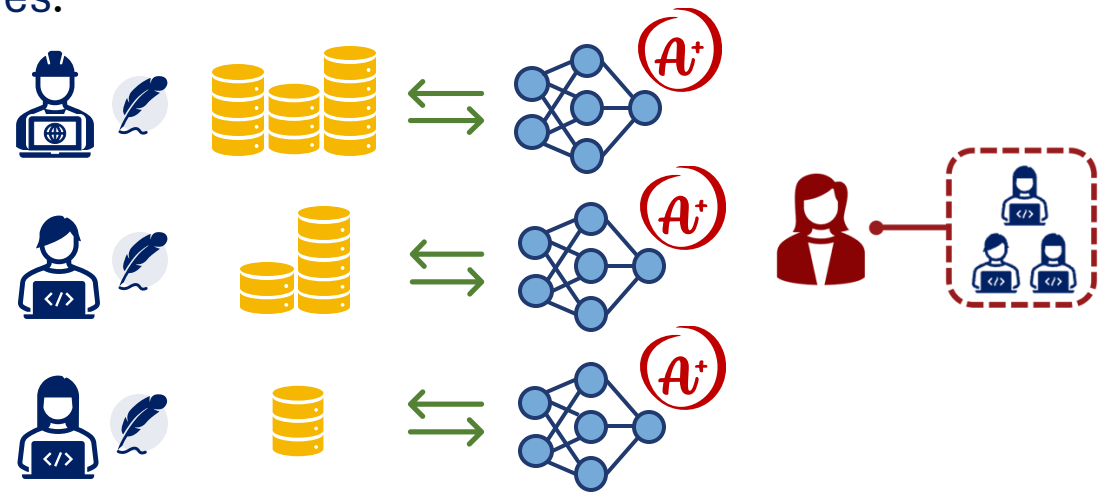

The incentive problem of CML

Conflict of interests + **Incomplete Information**

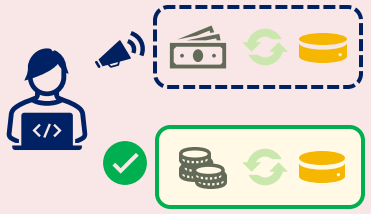

Contribution costs are often privately observable and difficult to verify.

Principal-agent problem > Monopolistic screening

Contract Design with Model Rewards

- Motivation
- **Models as Rewards**
- Optimal Contract
- Results

Model rewards are different from money:

(1) Non-rivalrous.

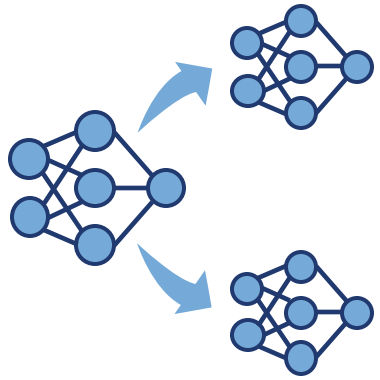

(2) Stochastic ex-ante.

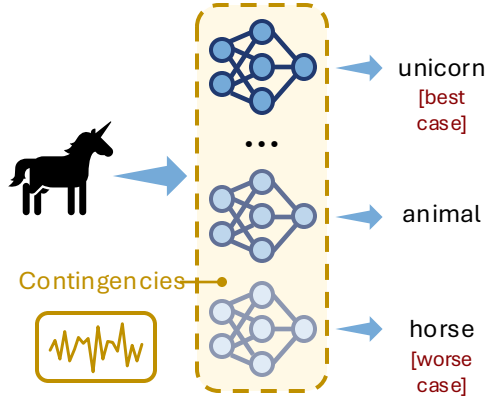

Assigning **deterministic model rewards** risks insolvency or under-utilization of the budget.

CML with Optimal Contract Design

- Motivation
- **Models as Rewards**
- Optimal Contract
- Results

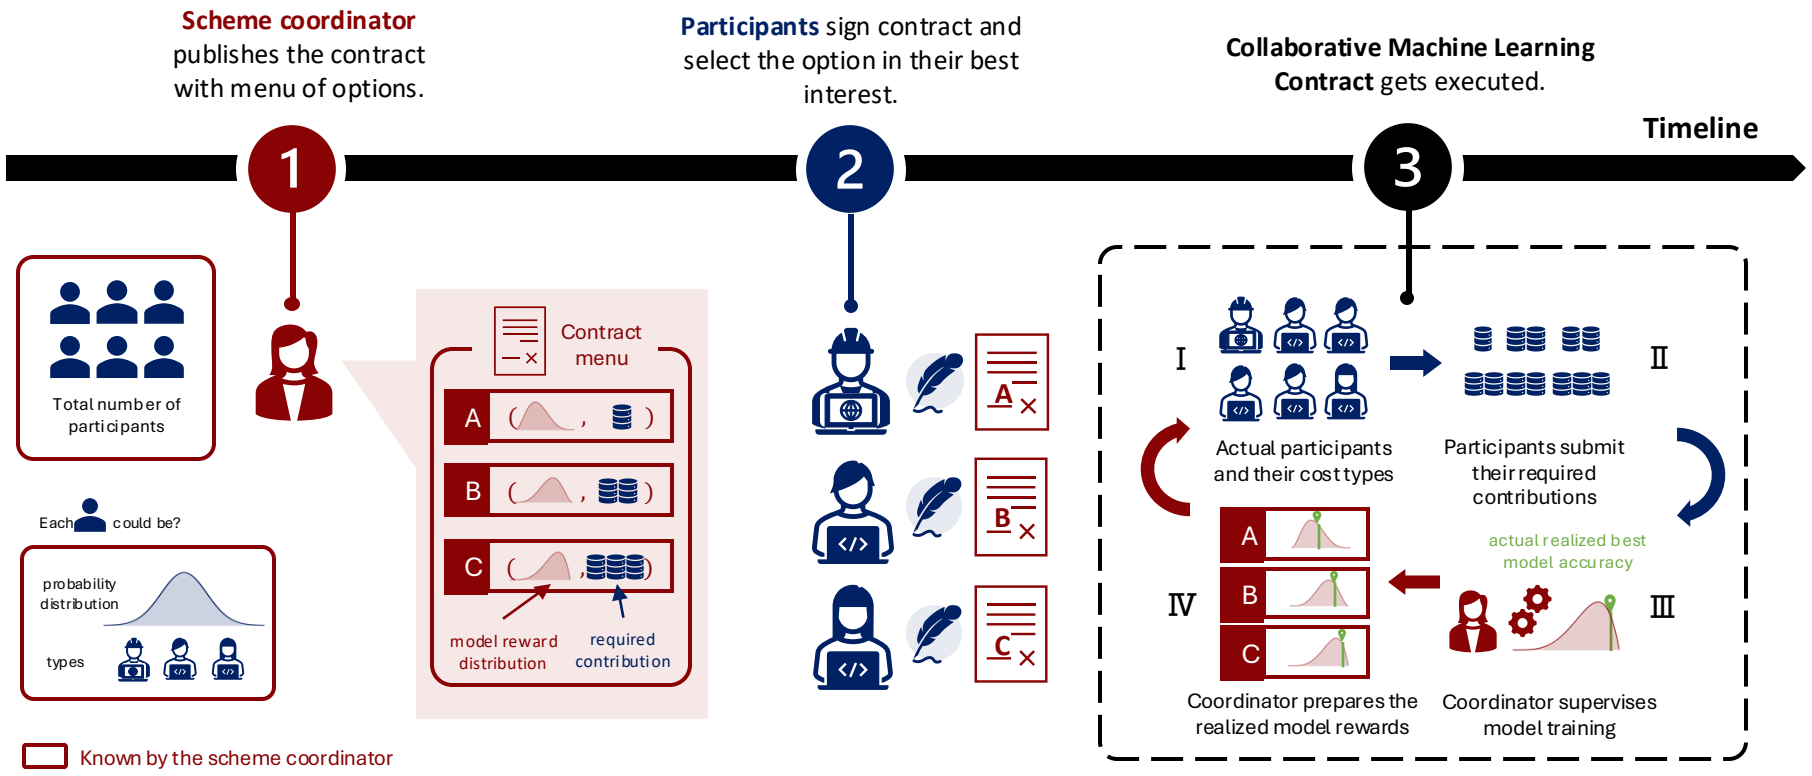

- Motivation
- Models as Rewards
- **Optimal Contract**
- Results

Solving for the Optimal Contract

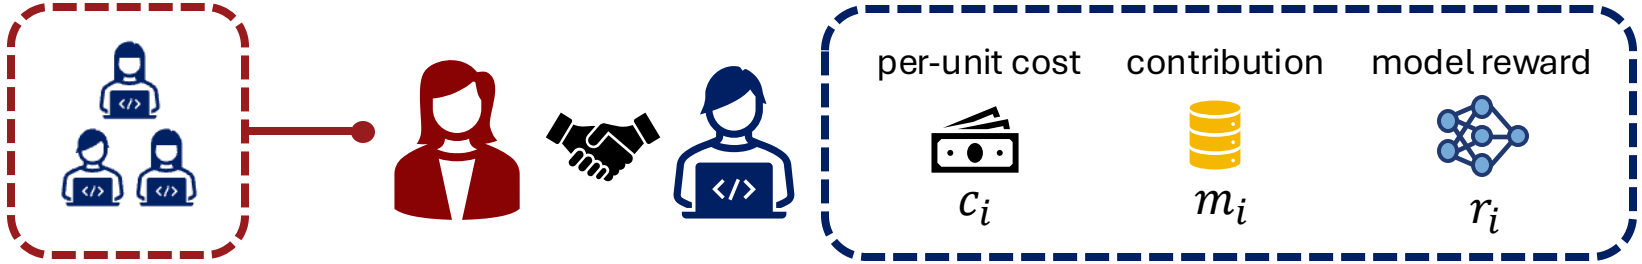

Coordinator's Utility Function

$$\mathbb{E}_{n\sim \text{Multi}(N,p)}\left[a\left(\sum_{i=1}^{I} n_i m_i\right)\right]$$

accuracy

Coordinator only knows the probability distribution.

Party's Utility Function

$$\mathbb{E}_{n_i \geq 1}[v(r_i)] - c_i m_i$$

revenue cost

Party knows they will be contributing when joining the scheme.

Constrained Optimization

$$\max_{(\mathbf{r}_i, m_i)_{i=1}^{I}} \mathbb{E}_{n \sim \text{Multi}(N,p)} \left[a \left(\sum_{i=1}^{I} n_i m_i \right) \right]$$

$$\text{s.t.} \begin{cases} \mathbb{E}_{n_i \geq 1}[v(r_i)] - c_i m_i \geq f_i, \forall i \\ \text{Joining the scheme gives each party a weakly higher net profit than opting out.} \end{cases}$$

Individual Rationality

Constrained Optimization

$$\max_{(\mathbf{r}_i, m_i)_{i=1}^{I}} \mathbb{E}_{n \sim \text{Multi}(N,p)} \left[a \left(\sum_{i=1}^{I} n_i m_i \right) \right]$$

$$\text{s.t.} \begin{cases} \mathbb{E}_{n_i \geq 1}[v(r_i)] - c_i m_i \geq f_i, \forall i \\ \mathbb{E}_{n_i \geq 1}[v(r_i)] - c_i m_i \geq \mathbb{E}_{n_j \geq 1}[v(r_j)] - c_i m_j, \forall i, j \end{cases}$$

Individual Rationality

Incentive Compatibility

Each party has incentive to tell the truth by choosing the option designed for them.

- Motivation
- Models as Rewards
- **Optimal Contract**
- Results

Constrained Optimization

$$\max_{(\mathbf{r}_i, m_i)_{i=1}^{I}} \mathbb{E}_{n \sim \text{Multi}(N,p)} \left[a \left(\sum_{i=1}^{I} n_i m_i \right) \right]$$

$$\text{s.t.} \begin{cases} \mathbb{E}_{n_i \geq 1}[v(r_i)] - c_i m_i \geq f_i, \forall i \\ \mathbb{E}_{n_i \geq 1}[v(r_i)] - c_i m_i \geq \mathbb{E}_{n_j \geq 1}[v(r_j)] - c_i m_j, \forall i, j \\ \|r(n)\|_\infty \leq a\left(\sum_{i=1}^{I} n_i m_i\right), \forall n \in \text{Multi}(N,p) \end{cases}$$

Individual Rationality

Incentive Compatibility

Budget Constraint

The administered model rewards must never surpass the collectively trained model in performance.

Hard to solve directly.

First-moment Problem

- Motivation
- Models as Rewards
- **Optimal Contract**
- Results

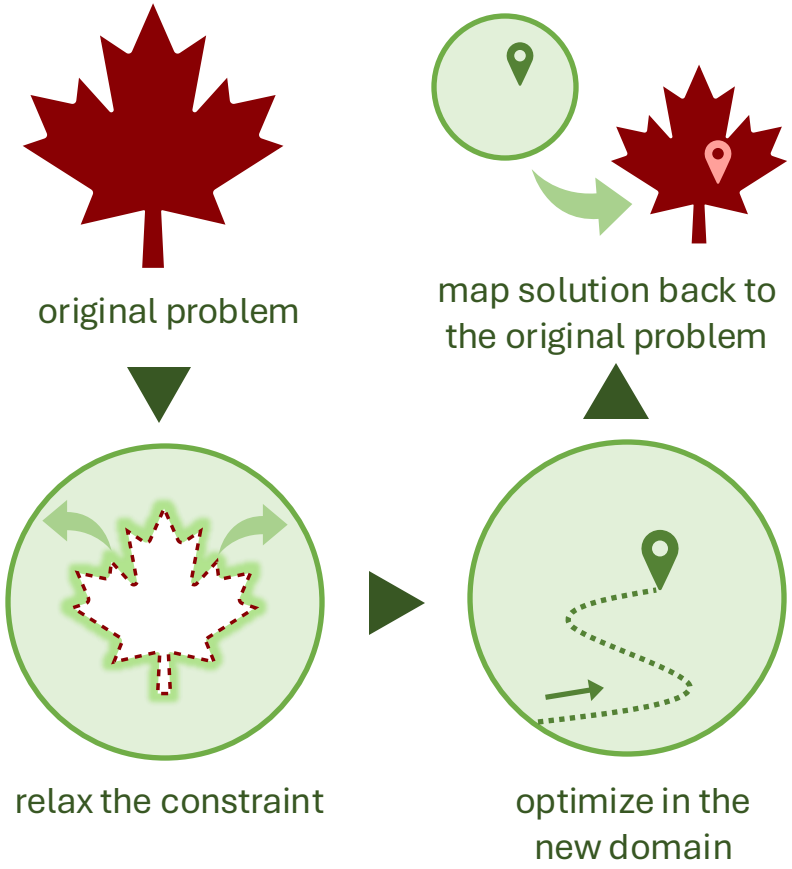

$$\max_{(t_i,m_i)_{i=1}^{I}} \mathbb{E}_{n\sim\text{Multi}(N,p)}\left[a\left(\sum_{i=1}^{I} n_i m_i\right)\right]$$

$$\text{s.t.}\begin{cases} t_i - c_i m_i \geq f_i, \forall i \\ t_i - c_i m_i \geq t_j - c_i m_j, \forall i, j \\ t_i \leq \mathbb{E}_{n_i \geq 1}\left[v\left(a\left(\sum_{i=1}^{I} n_i m_i\right)\right)\right], \forall i \end{cases}$$

- ☑ Fewer variables
- ☑ Fewer constraints
- ☑ Convex

Experiment Results

- Motivation
- Models as Rewards
- Optimal Contract
- **Results**

training a model is manageable
$f_i > 0$

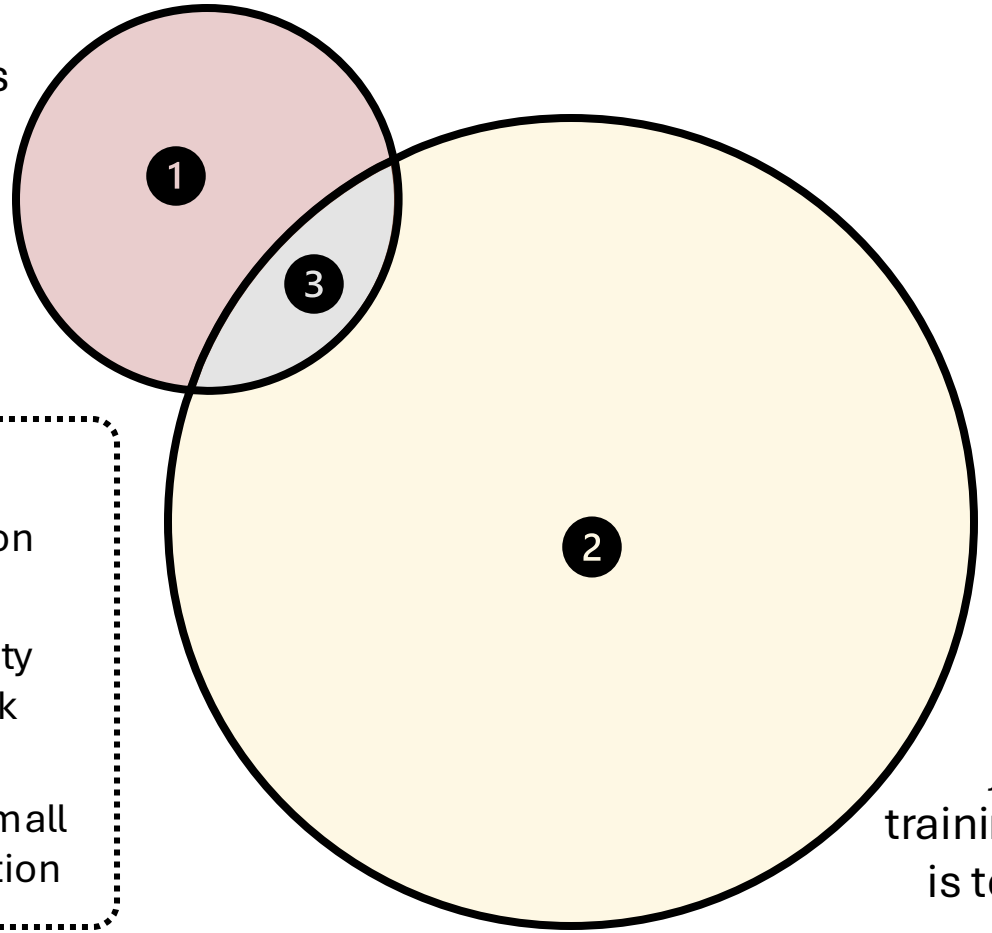

$f_i = 0$
training a model is too costly

1 Big-firm cooperation

2 Small-party teamwork

3 Big-and-small collaboration

Experiment Results (Big-firm cooperation)

- Motivation
- Models as Rewards
- Optimal Contract
- **Results**

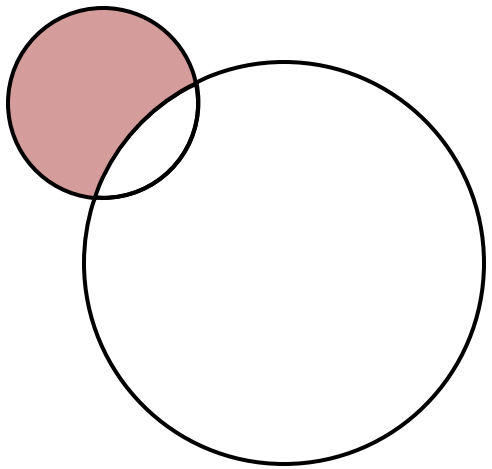

Data contributions

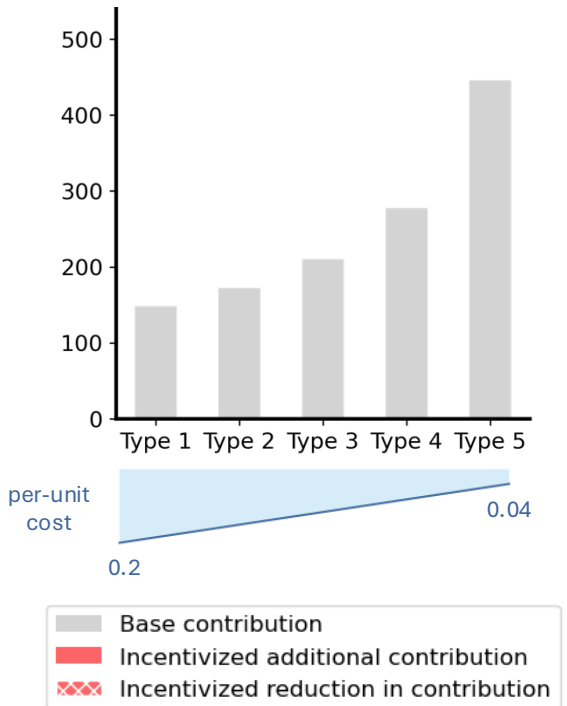

Model rewards

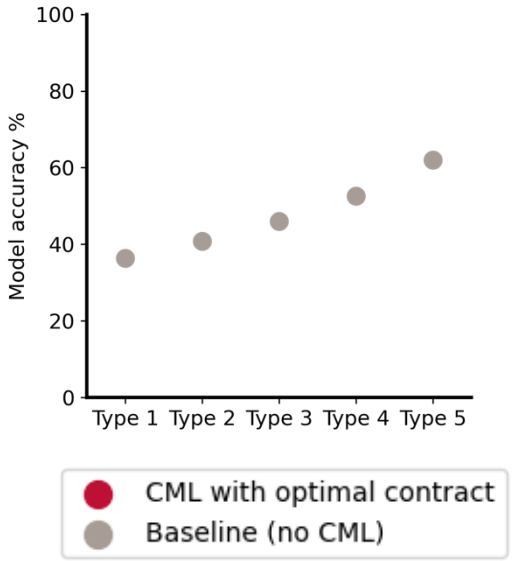

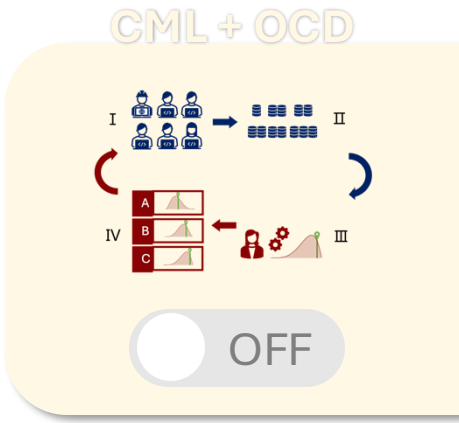

Experiment Results (Big-firm cooperation)

- Motivation
- Models as Rewards
- Optimal Contract
- **Results**

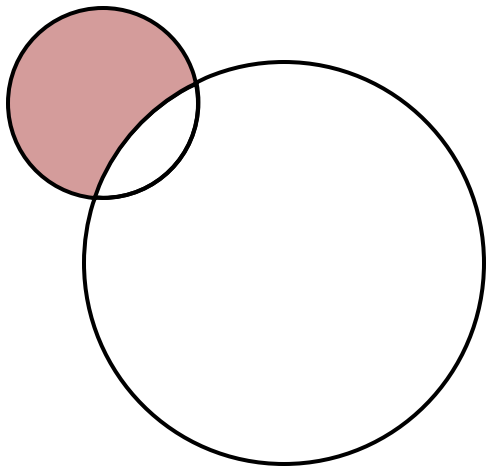

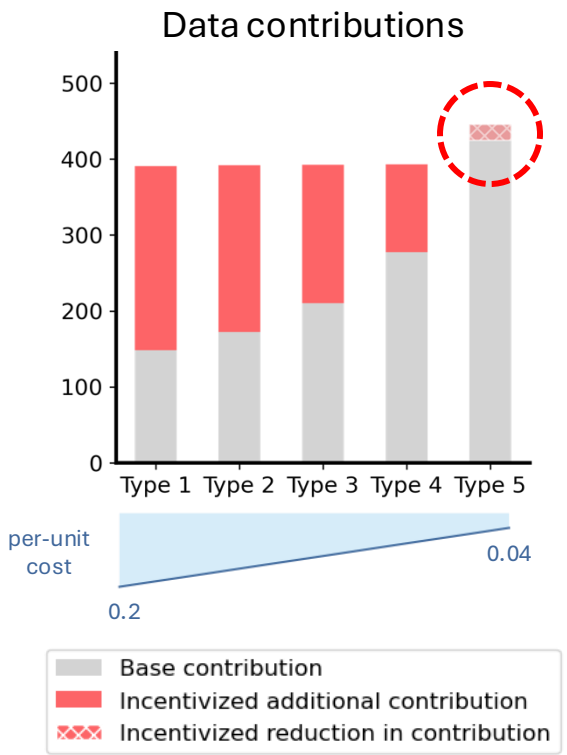

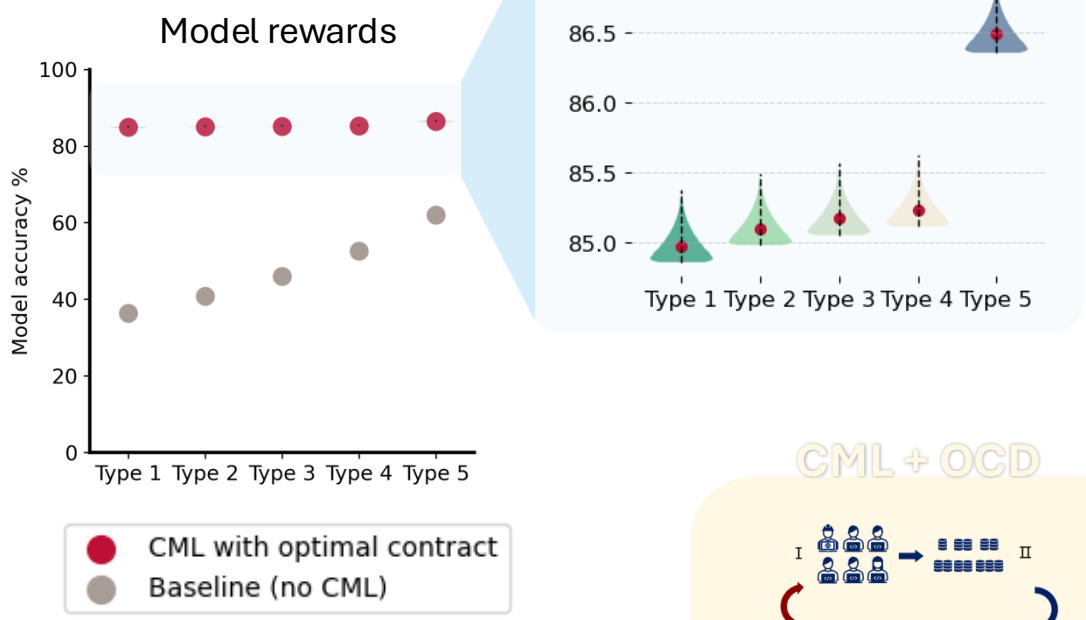

A party can be incentivized to contribute less.

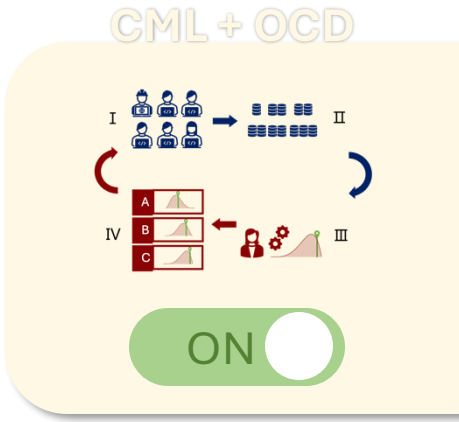

Experiment Results (Small-party teamwork)

- Motivation
- Models as Rewards
- Optimal Contract
- **Results**

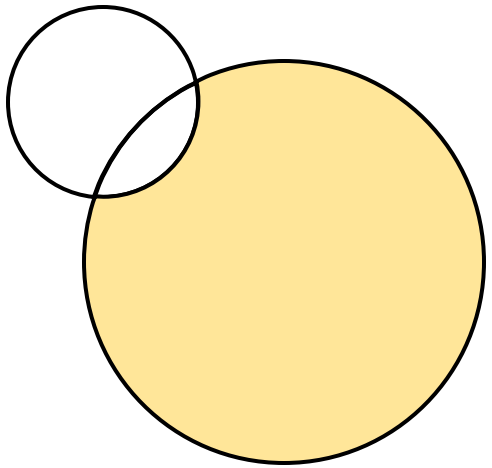

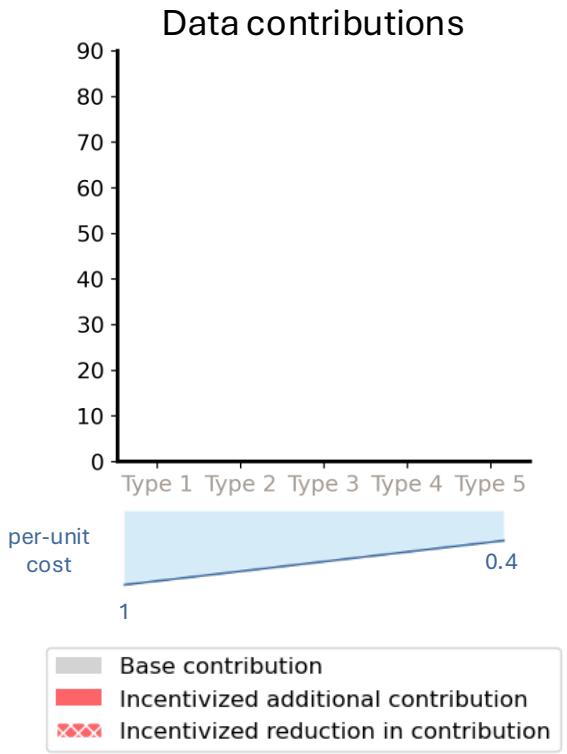

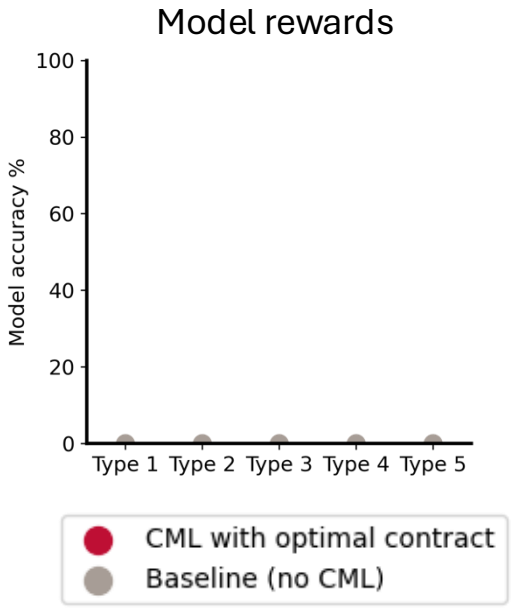

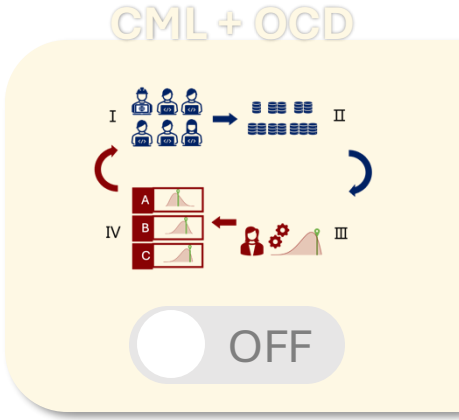

Experiment Results (Small-party teamwork)

- Motivation
- Models as Rewards
- Optimal Contract
- **Results**

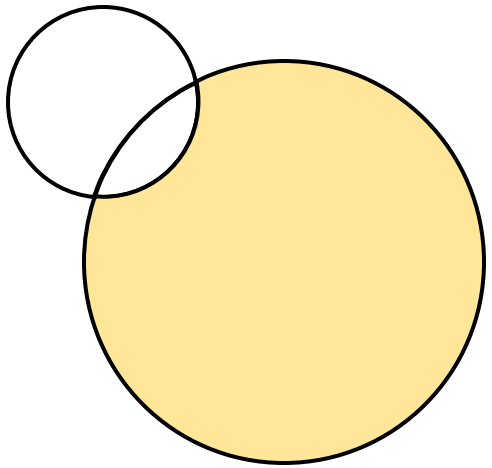

Data contributions

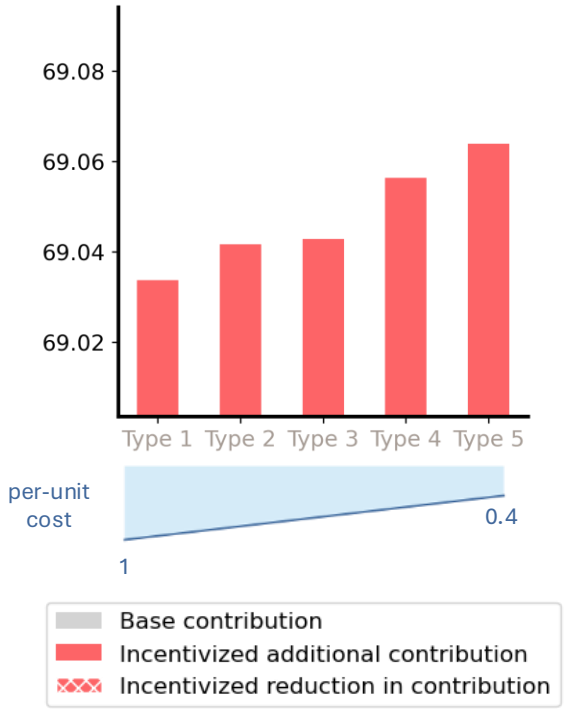

Model rewards

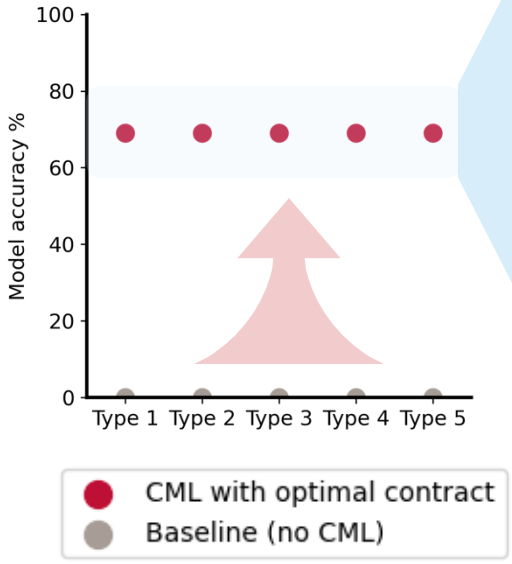

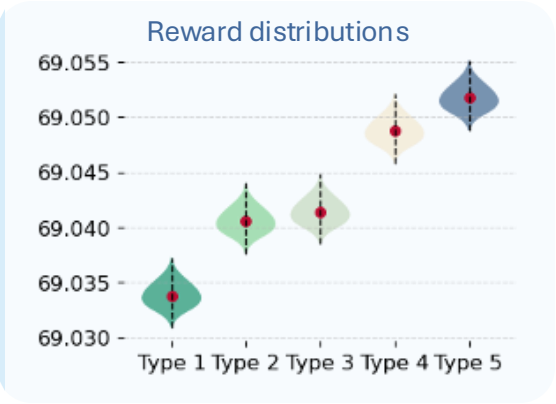

Parties successfully surmount model training barrier.

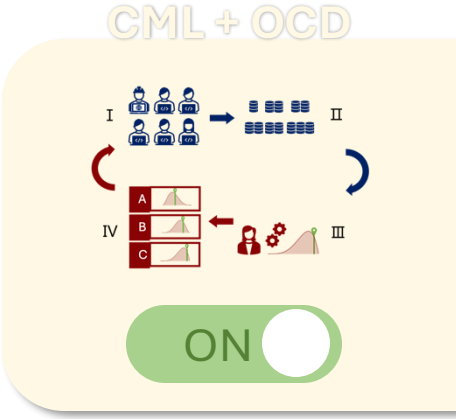

Experiment Results (Big-and-small collaboration)

- Motivation
- Models as Rewards
- Optimal Contract
- **Results**

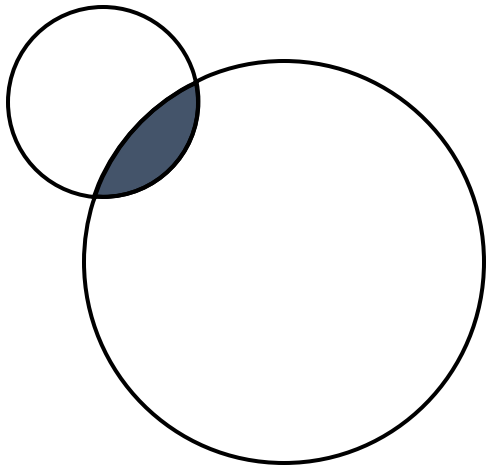

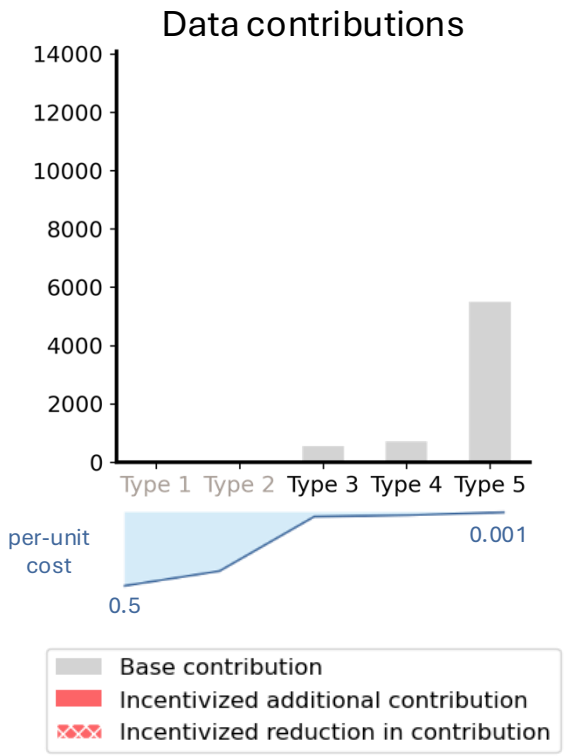

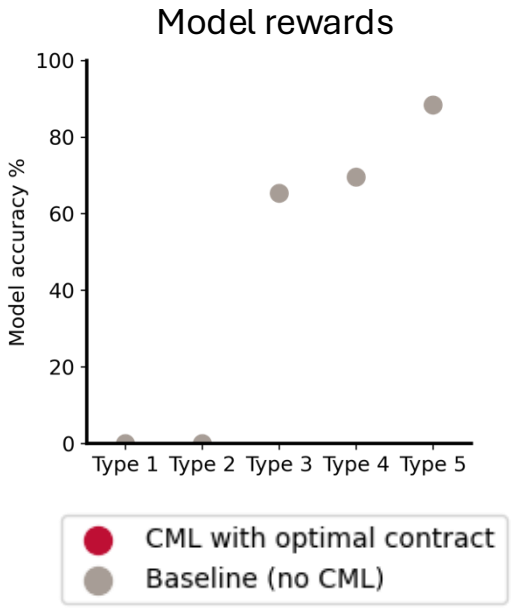

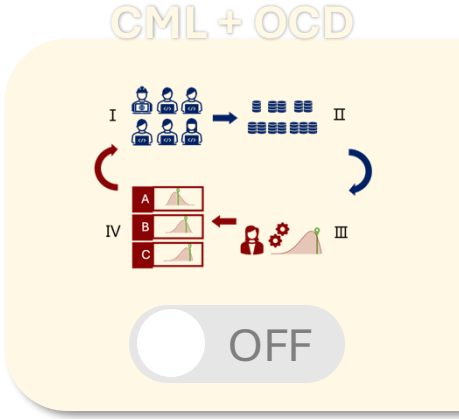

Experiment Results (Big-and-small collaboration)

- Motivation
- Models as Rewards
- Optimal Contract
- **Results**

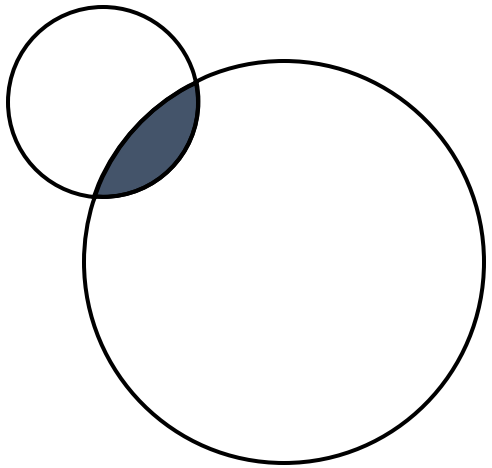

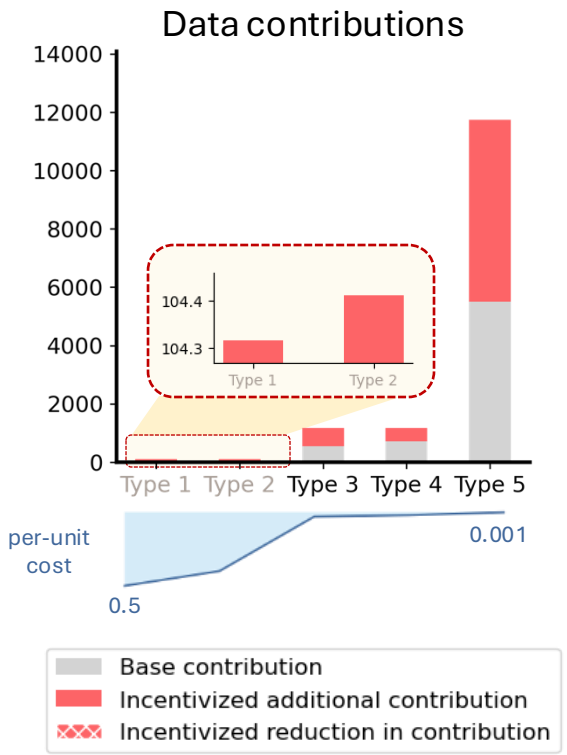

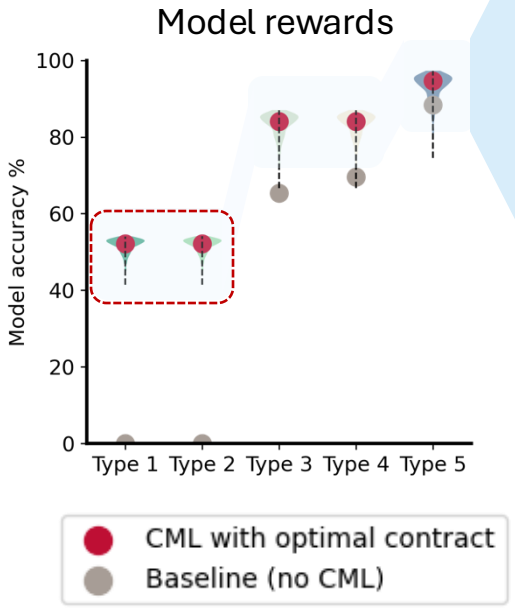

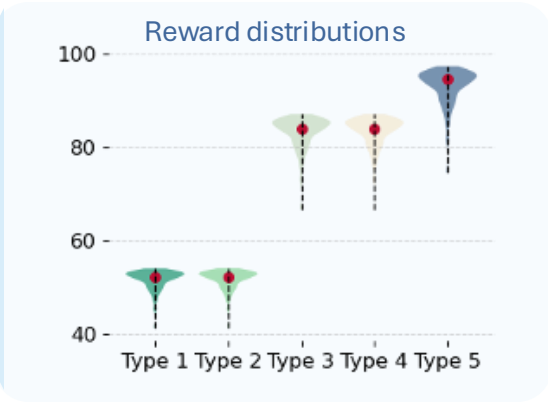

Small parties can still gain from collaboration.

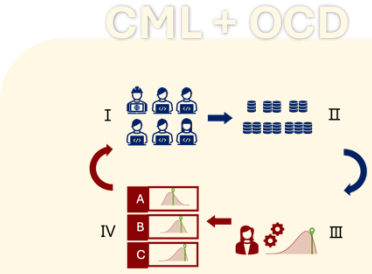

Conclusion

"You do not rise to the level of your goals. You fall to the level of your systems."
— James Clear, *Atomic Habits*.

A promising crowdsourcing paradigm to democratize AI?

+

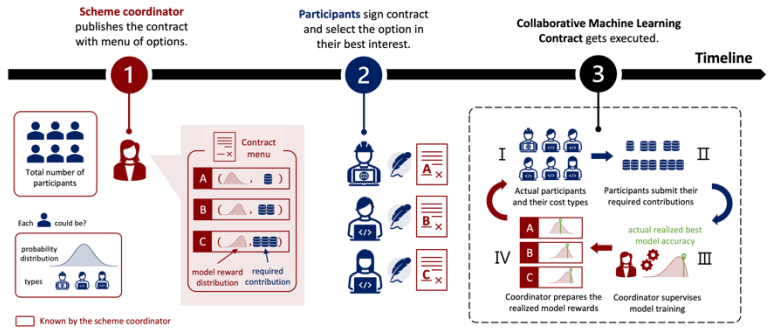

Full paper

Code

Limitations & Future Work

Full paper

Code

Empirical justification for behavioral assumptions

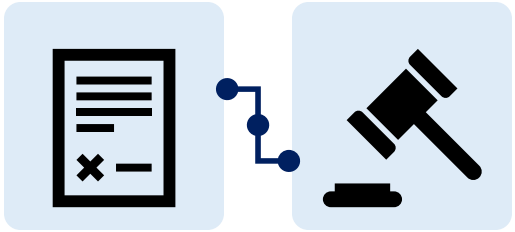

contract design — auction design

Prior-training accuracy function & valuation function

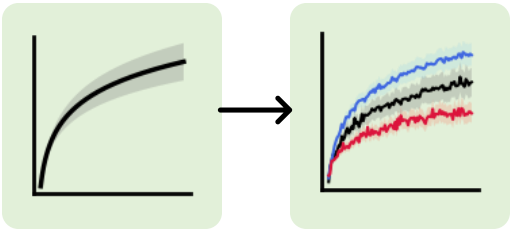

consensus — pluralism

Incorporating quality in the contribution measure

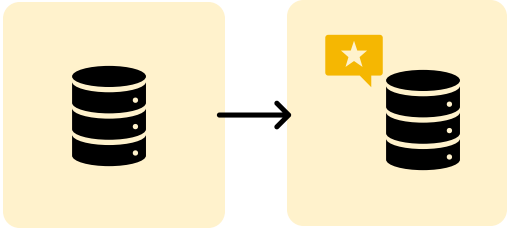

homogeneous contribution — heterogeneous contribution

Distributional assumption & combinatorial challenge

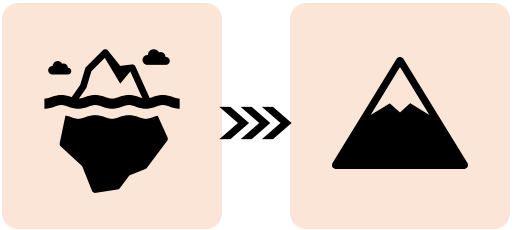

multinomial distribution — effective approximation