Multi-Objective Learning:
A Unifying Framework for Collaboration, Fairness, and Robustness

Nika Haghtalab

EECS, UC Berkeley

More Data ... More Stakeholders

Learning when data comes from many populations and agents and learning guarantees must meet the needs of many populations and agents

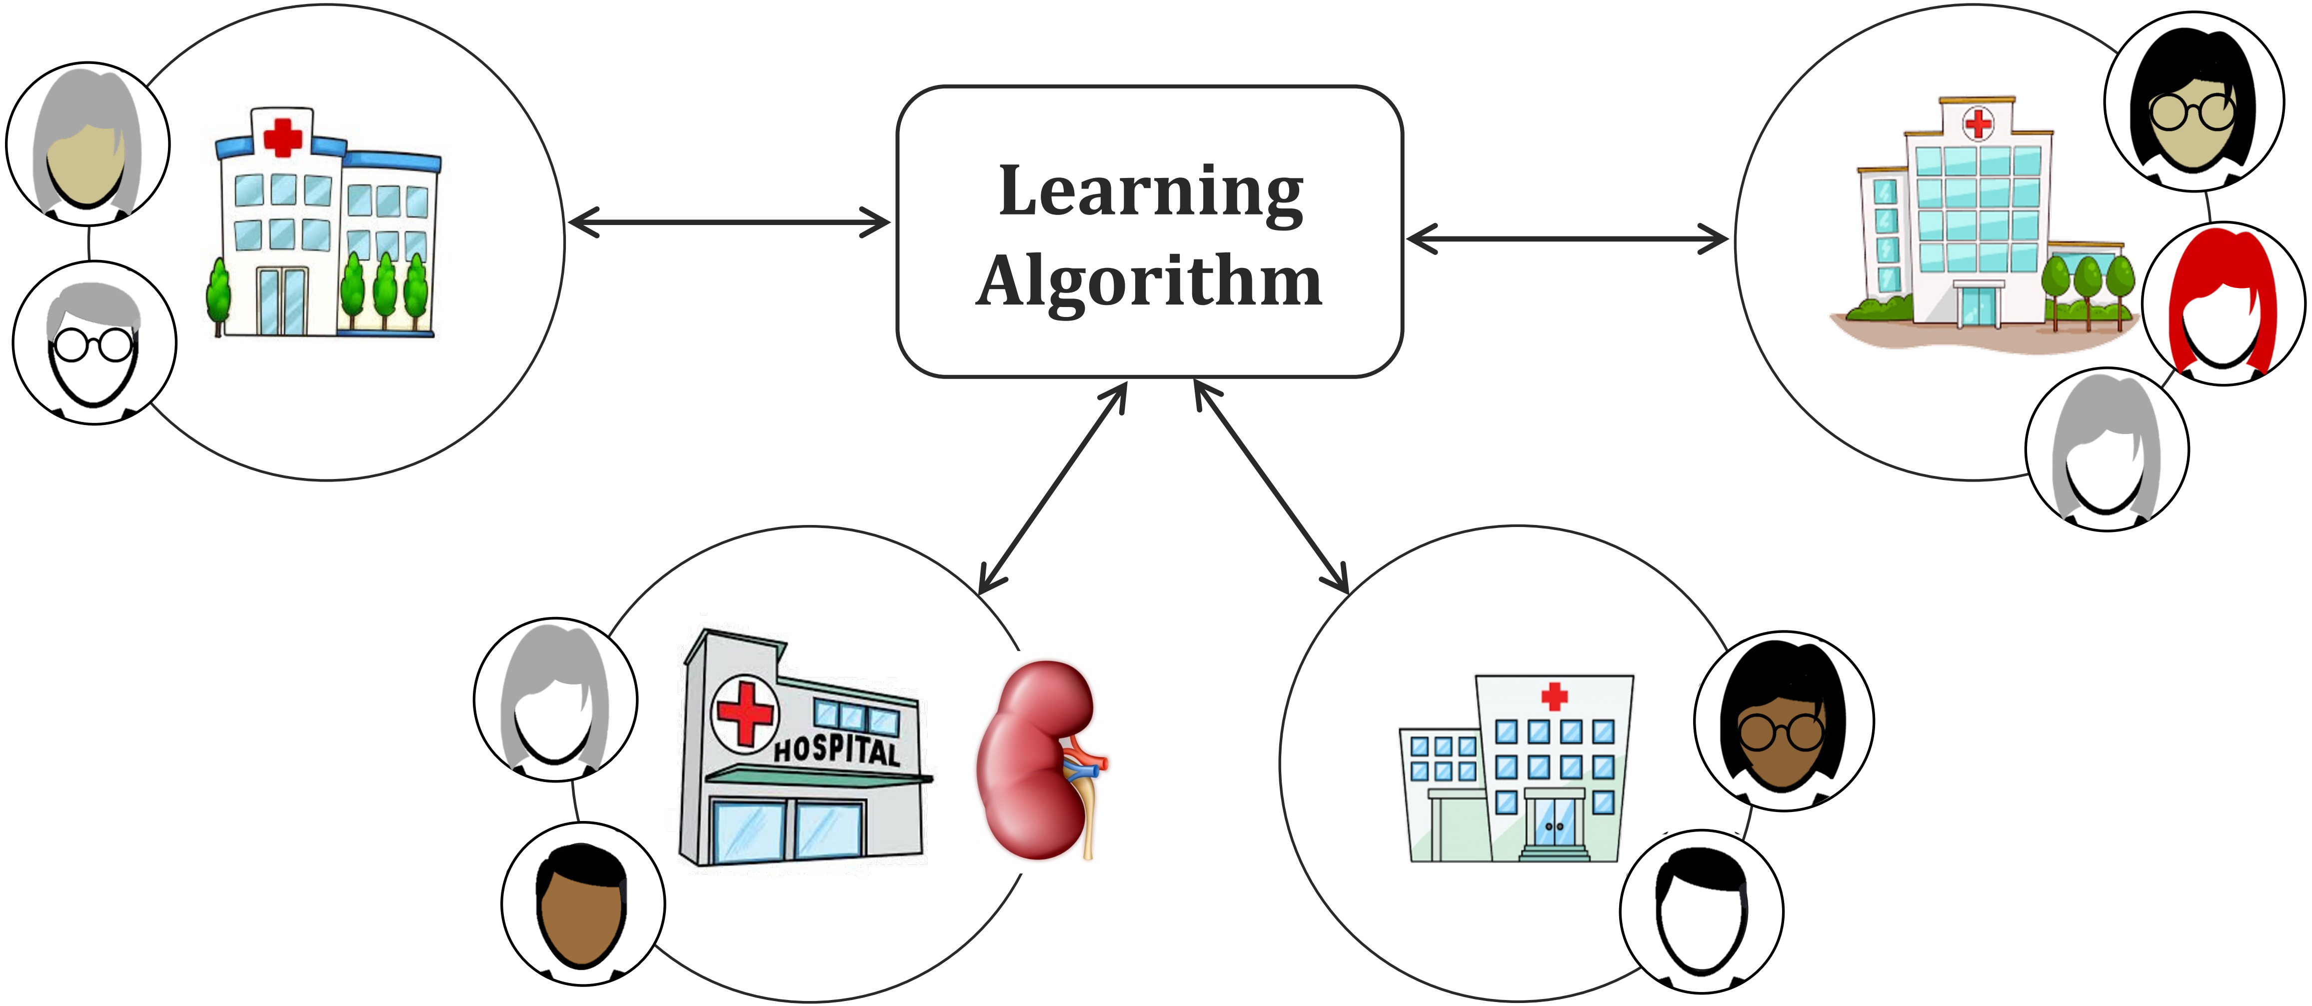

More Data … More Stakeholders

Learning when data comes from many populations and agents and learning guarantees must meet the needs of many populations and agents

Data sharing: E.g., large genome studies require large-scale data-sharing and collaboration between many institutes and research labs.

Per-group Guarantees: E.g., medical research must provide solutions that apply to different localities, populations, threat models, and resources.

Cost-Benefit Tradeoffs: E.g., taking samples in physical domains is costly to individuals and data curators, e.g., medical tests, lead pipe testing, …

The Issue with On-Average Guarantees

Typical learning algorithms work well **on average** over the data sources

- Good for when the data is homogenous across sources
- Good for learning across data centers.

Human and organization data: highly heterogenous.

Learning difficulty and exhibited patterns vary significantly
→ Some populations are easier to learn than others.
→ On-average guarantees don't lead to meaningful solutions for all subpopulations.

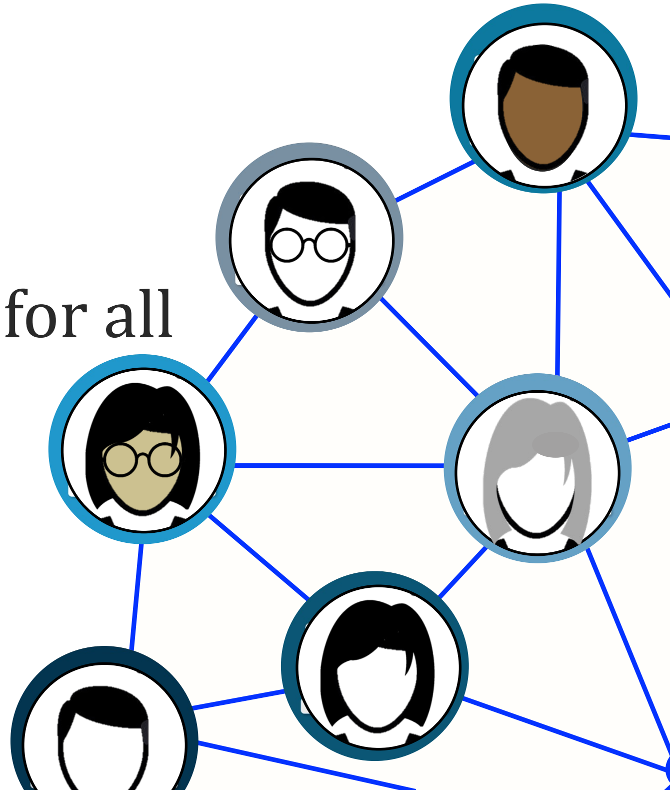

Example 1: Calibrated (valid) predictions

A predictor $p: X \to [0,1]$ is *calibrated* if for all v, $\mathbb{E}_{(x,y)}[y|p(x) \approx v] \approx v$.

[Dawid'82]

E.g., if p predicts heart failure probability in a patient, then among patients where p predicts 0.1, 10% truly develop heart failure.

Classical Calibration

Rather a weak guarantee.

Met by $p(x) := \mathbb{E}[y]$.

Not meaningful for any individual.

[Dawid'82]

Per-Group Guarantee

Holding for each set of features in class $\mathcal{S} \subseteq 2^X$

all $v, S \in \mathcal{S}$

$$\mathbb{E}[y|p(x) \approx v, x \in S] \approx v.$$

When $\mathcal{S} = 2^X$, satisfies $p(x) \approx \mathbb{E}[y|x]$

Formalized by Herbert-Johnson Kim Reingold, Rothblum '18, also Foster Kakade'06, Sandroni Smorodinsky Vohra'03

Importance of Calibration across groups

Predictors are used for deciding treatment plans.

Miscalibration across groups leads to different quality of care.
E.g., Empirical study of miscalibration of medical risk predictors across race. [Obermeyer et al '19]

Multi-calibration:

→ Intuitively, predictions mean the same across group.

→ Calibration alone can stereotype entire populations.

  → $p(x) := \mathbb{E}[y | x \in S]$

→ The more intersectional groups considered in multi-calibration, the closer a multi-calibrated predictor is to Bayes probability.

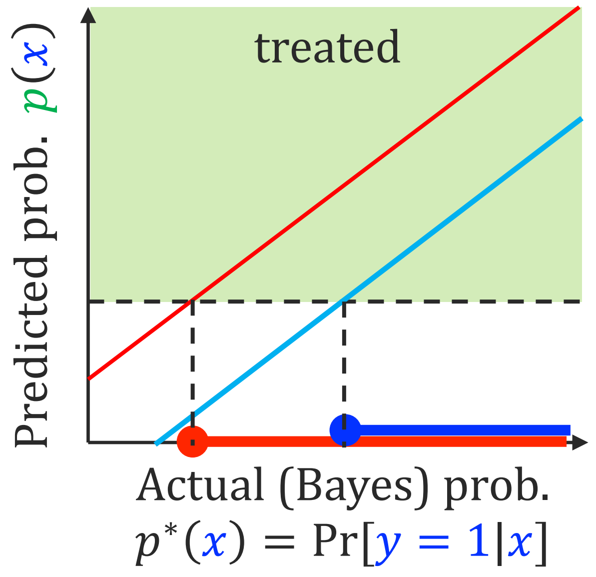

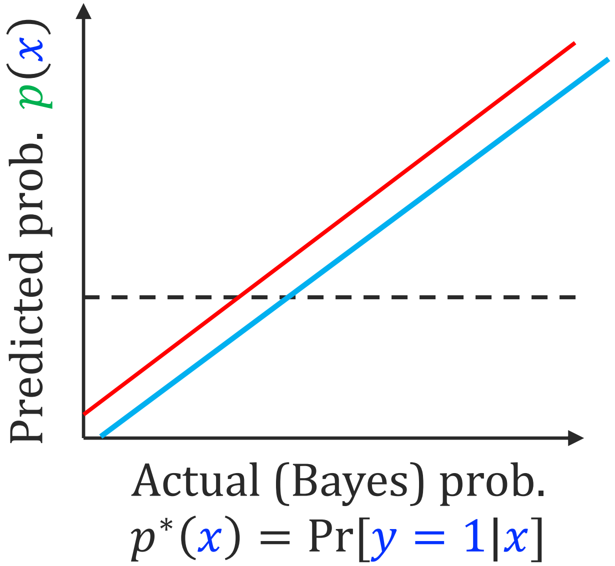

Example 2: Accuracy / Prediction Loss

A function $h: X \to [0,1]$ has loss $L_D(h) = \mathbb{E}_{(x,y)}[\ell(y, h(x))]$.

E.g., binary classification, regression losses.

Classical PAC

Learn h with $L_D(h) \le \epsilon$.
Even if $L_D(h) \le 0.05$, h could have **50% error for $^1/_{10}$ of the population.**

Problematic when that $^1/_{10}$ of the population correlates with a type.

[e.g. Valiant '84]

Per-Group Guarantee

Holding for each one of given distributions $\mathcal{D} = \{D_1, \dots, D_k\}$
all $D_i \in \mathcal{D}$, $L_{D_i}(h) \le \epsilon$.

[formalized by Blum **H.** Procaccia Qiao'17. Related formalisms Kearns Neel Roth Wu'18, Mohri Sivek Suresh '19, Sagawa Koh Hashimoto Liang'20, also literature on domain adaptation.]

Example 3: Accuracy vs. Cost Tradeoffs

A function $h: X \to [0,1]$ has loss $L_D(h) = \mathbb{E}_{(x,y)}[\ell(y, h(x))]$.

E.g., binary classification, regression losses.

Classical Tradeoffs

Optimal total number of samples to collectively guarantee

all $D_i \in \mathcal{D},\ L_{D_i}(h) \le \epsilon$.

Unfairly distribute the burden over distributions.

Per-Group Tradeoffs

Guaranteeing for all $D_i \in \mathcal{D},\ L_{D_i}(h) \le \epsilon$ with *optimal per-distribution* number of samples.

Ensuring optimal tradeoffs per distribution.

[formalized by Blum **H.** Phillips, Shao '21]

Learning Across Multiple Objectives

Enabling learning processes that satisfy **multiple objectives**

for **several agents** from **collectively fewer resources**.

Practical Applications and considerations

- **Data sharing and collaborative learning:** In use across networks of devices, hospitals, etc., and behind recent major scientific discoveries.
- **Robustness:** Successful deployment in any one of possible scenarios.
- **Fairness:** Ensuring welfare of representative subpopulations.

Foundations of
Multi-Objective Statistical Learning

Rest of this talk:

Multi-Objective Learning, one unifying framework

1. Collaborative and multi-distribution learning
2. Multi-calibration
3. Optimal tradeoffs

Multi-Objective Learning: Per-Group Guarantees

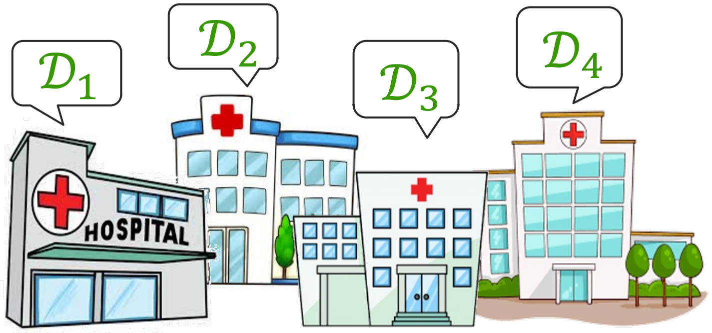

There are k populations (distributions), represented by unknown $\mathcal{D}_1, \mathcal{D}_2, \ldots, \mathcal{D}_k$ from which we can sample.

There are r possible loss functions $\ell^1, \ldots, \ell^r$, e.g., $\ell^j(x, y, f) = 1(y \neq f(x))$ or $\left(y - f(x)\right)1(f(x) = v)$, and $\mathrm{L}^j_{\mathcal{D}_i}(f) = \mathbb{E}_{(x,y)\sim\mathcal{D}_i}\left[\ell^j(x, y, f)\right]$

Multi-Objective Learning

Learn a function f that is simultaneously good for every population and every loss function in consideration.

$$\max_{j\in[r]} \max_{i\in[k]} \mathrm{L}^j_{\mathcal{D}_i}(f) \leq \epsilon \qquad \text{(uncovering a universally good model)}$$

$$\max_{j\in[r]} \max_{i\in[k]} \mathrm{L}^j_{\mathcal{D}_i}(f) - \min_{h^*\in H} \max_{j\in[r]} \max_{i\in[k]} \mathrm{L}^j_{\mathcal{D}_i}(h^*) \leq \epsilon \qquad \text{(best in class } H\text{)}$$

Multi-Calibration and Multi-Distribution Learning

Multi-Calibration

Consider the class of all predictor $H = [0,1]^X$, the set of loss functions $\ell^{v,S,\sigma}(x, y, p) = \sigma\big(y - p(x)\big)1(p(x) \approx v)\ 1(x \in S)$ for $\sigma \in \{-1, +1\}$, predicted values v and subgroups $S \in \mathcal{S}$, and a single distribution $\mathcal{D}$.

p is $(\mathcal{S}, \epsilon)$-multicalibrated if

$$\max_{\sigma,v,S} \mathrm{L}_{\mathcal{D}}^{\sigma,v,S}(p) - \overbrace{\min_{p^* \in H} \max_{\sigma,v,S} \mathrm{L}_{\mathcal{D}}^{\sigma,v,S}(p^*)}^{=\,0 \text{ by Bayes classifier}} \leq \epsilon$$

Collaborative/Multi-Distribution Learning

For any hypothesis class H and any k distributions $\mathcal{D}_1, \ldots, \mathcal{D}_k$, learn function f such that

$$\max_{i \in [k]} \mathrm{L}_{\mathcal{D}_i}(f) - \min_{h^* \in H} \max_{i \in [k]} \mathrm{L}_{\mathcal{D}_i}(h^*) \leq \epsilon$$

A Unifying Perspective

In the past 5-7 years, similar models were introduced by several different communities. Mostly inspired by ideas of fairness, robustness, and collaborations.

→ Collaborative Learning [Blum, **H**, Procaccia, Qiao '17]

→ $\mathcal{D}_i$s represent agent distributions. Agents are willing to collaborate.

→ Agnostic (Fair) Federated Learning [Mohri, Sivek, Suresh'19]

→ $\mathcal{D}_i$s represent client distributions. Fairness goals and implications.

→ (Group) Distributionally Robust optimization [Sagawa, Koh, Hashimoto, Liang '19]

→ $\mathcal{D}_i$s represent possible distribution shifts. Robustness and fairness goals.

→ Multi-group Agnostic PAC [Rothblum, Yona'21]

→ $\mathcal{D}_i$s represent subpopulations and loss functions capture regret to optimal loss

→ Muti-Calibration [Herbert-Johnson, Kim, Reingold, Rothblum '18]

→ $\mathcal{D}$ a single distribution, with calibration loss functions $\big(y - f(x)\big)1(f(x) = v, x \in S)$ taking membership and predicted value.

→ Multi-group Fairness [Kearns, Neel, Roth, Wu'18]

→ $\mathcal{D}$ a single distribution, loss functions capturing errors of various types on $1(x \in S)$.

Challenges of meeting multiple objectives

Multi-Objective Learning

Learn a function f that is simultaneously good for every population and every loss function in consideration.

$$\max_{j\in[r]} \max_{i\in[k]} \mathrm{L}^{j}_{\mathcal{D}_i}(f) \leq \epsilon$$

(uncovering a universally good model)

Algorithms with on-average guarantees perform poorly on groups, even on simple datasets.

→ Some distributions are easier to learn than others.

→ Also depends on similarity across different populations.

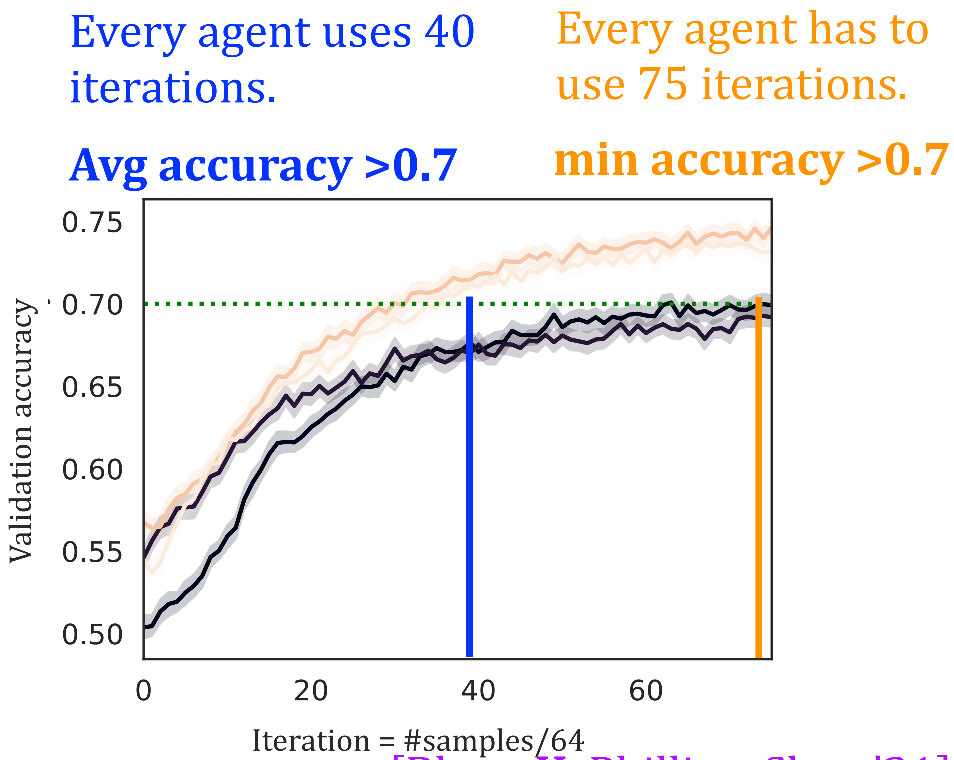

[Blum, H, Phillips, Shao ’21]

Can we provide multi-objective learning guarantees from reasonably small amount of data?

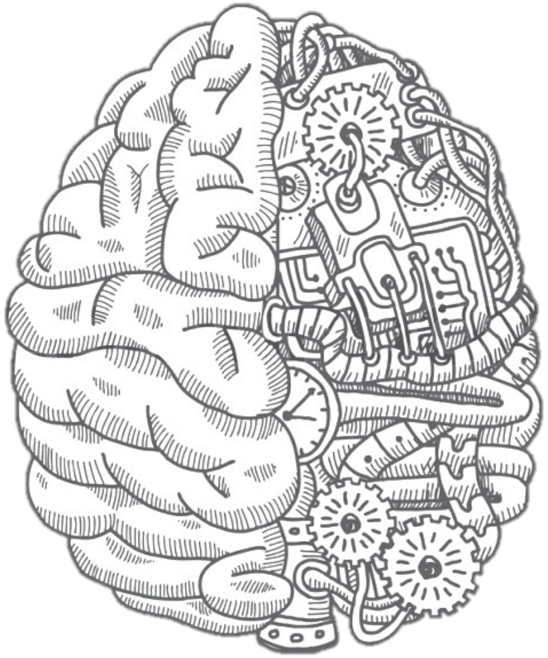

From One to Multiple Distributions (and Objectives)

Well-developed theory for how much resources are needed to learn a single distribution under one loss.

Import insights, algorithms, techniques, etc., from the single distribution setting.

For comparison

One Distribution (Statistical Learning)

Given sample access to an unknown $\mathcal{D}$,

find f, s.t. with high probability,

$$\mathrm{L}_{\mathcal{D}}(f) \leq \min_{h^* \in H} \mathrm{L}_{\mathcal{D}}(h^*) + \epsilon$$

Multiple Distributions

Given sample access to unknown $\mathcal{D}_1, \ldots, \mathcal{D}_k$,

find f, s.t. with high probability,

$$\max_{i \in [k]} \mathrm{L}_{\mathcal{D}_i}(f) \leq \min_{h^* \in H} \max_{i \in [k]} \mathrm{L}_{\mathcal{D}_i}(h^*) + \epsilon$$

ERM's Sample complexity: $\tilde{\Theta}\left(\frac{\boxed{VCD(H)}}{\epsilon^2}\right)$

Proxy for # parameters

Algorithm?

Sample complexity?

Foundations of
Multi-Objective Statistical Learning

Rest of this talk:

- Multi-Objective Learning, one unifying framework

1. Collaborative and multi-distribution learning
2. Multi-calibration
3. Optimal tradeoffs

Multi-Distribution Learning Needs Interactions

Standard algorithms and settings lack interactions

- # of samples, learning rates, and update frequencies decided non-interactively.
- Ignores varying distribution difficulty and relevance.

Non-Interactive

Sample complexity of existing algorithms, for k agents/dists $= \Theta(k) \times$ Learning for 1 distribution separately

[Blum, **H**, Procaccia, Qiao '17]

Without an "interactive" protocol,
Multi-Distribution learning (almost) ineffective at saving resources across multiple learning tasks.

Problematic for fairness and multi-agent collaboration purposes.

Multi-Objective Learning Needs Interactions

Standard algorithms and settings lack of interactions

- # of samples, learning rates, and update frequencies decided non-interactively.
- Ignores varying distribution difficulty and relevance.

Non-Interactive

Sample complexity of existing algorithms, for k agents/dists $= O\left(k \ln(k) \cdot \frac{\log(|H|)}{\epsilon^2}\right)$

In this regime
Group DRO
Multi group agnostic
Agnostic Federated Learning
...

Interactivity

To benefit from cross-learning, the distributions need to interact adaptively.

→ Decisions about $\mathcal{D}_i$ must depend on how well $\mathcal{D}_i$ has done so far, compared to $\mathcal{D}_j$.

<u>Adjusting sample collection based on past performance</u>

There is an algorithm
Overall # samples $= O\left(\frac{\log(|H|)}{\epsilon^2} + \frac{k \ln(k)}{\epsilon^2}\right)$

[Blum, **H**, Procaccia, Qiao ’17]

[**H**, Jordan, Zhao ’22]

Interactive Protocol as Game Solving

Re-imagining multi-objective learning as a zero–sum game.

Approximate MinMax equilibrium

$$\max_{i\in[k]} \mathrm{L}_{\mathcal{D}_i}(f) \le \min_{h^*\in H} \max_{i\in[k]} \mathrm{L}_{\mathcal{D}_i}(h^*) + \epsilon$$

Minimizing Agent: 1
Minimize the loss over function class H

2 **Maximizing Agent:** Maximize the loss over the class of distributions $\mathcal{D}_1, \dots, \mathcal{D}_k$.

Using no-regret algorithms to find an approximate minmax equilibrium.

- Sufficient for one player to play no-regret, and another to best respond or be no-regret.
  → Best-Response v. No-regret, No-regret v. Best-Response, No-regret v. No-regret

- **Implementation considerations:**
  → $\mathrm{L}_{\mathcal{D}_i}(f)$ is estimated through sampling from $\mathcal{D}_i$, want quick convergence.

Interactive Dynamics and Algorithm Design

A player is best-response if her choice is per-step near optimal (or good enough).

A player is no-regret: if she received a loss(utility) that is near-optimal, in hindsight, over any sequence of actions played by the other player.

Min-player No-Regret

Play $f_1, \cdots f_T$ online and see adversarial choices $(\mathcal{D}_1, L^1), \ldots, (\mathcal{D}_\mathrm{T}, L^T)$, s.t.

$$\sum L^t_{\mathcal{D}_t}(f_t) \le \min_{h^*} \sum L^t_{\mathcal{D}_t}(h^*) + o(T)$$

Max-player No-Regret

Play $(\mathcal{D}_1, L^1), \ldots, (\mathcal{D}_\mathrm{T}, L^T)$ online and see adversarial choices $f_1, f_2, \cdots f_T$, s.t.

$$\sum L^t_{\mathcal{D}_t}(f_t) \ge \max_{\mathcal{D}^*, L^*} \sum L^*_{\mathcal{D}^*}(f_t) - o(T)$$

No-regret algorithms exists, if $L^j_{\mathcal{D}}$ can be accessed, with regret $\sqrt{T \ln(\#player\ actions)}$

Time-averaged actions' convergence rate to minmax equilibrium $=$ Min-Player Regret $+$ Max-Player Regret

[e.g., Freund & Schapire '96]

Dynamics 1: Best-Response versus No-Regret

An approach: Solve with a **no-regret algorithm** against a **best-responding agent**.

Min Player: The best-responding agent. For any distribution over [k], $\alpha_1^t, \ldots, \alpha_k^t$, it uses an Empirical Risk Minimizer to learn $h^t \in H$ on the distribution $P^t = \sum \alpha_i^t D_i$

Sample : **proportional to $\boldsymbol{\alpha_i^t}$.**

Max Player: The no-regret learning agent. Maintains a distribution over $[k]$, say weights $\alpha_1^t, \ldots, \alpha_k^t$ over the agents. Proxy of how poorly they've been doing so far.

Depending on $\boldsymbol{h^1, \ldots, h^{t-1}}$.

Sample

Why does this work?

Simplifying assumption: $\min_{h^* \in H} \max_{i \in [k]} \mathrm{L}_{\mathcal{D}_i}(h^*) = 0.$

Min Player: The best-responding agent. For any distribution over $[k]$, $\alpha_1^t, \dots, \alpha_k^t$, it uses an Empirical Risk Minimizer to learn $h^t \in H$ on the distribution $P^t = \sum \alpha_i^t D_i$.

$$\mathrm{L}_{P^t}(h^t) \le \epsilon' \qquad \text{Samples } O\left(\frac{\log(|H|)}{\epsilon'^2}\right)$$

Max Player: The no-regret learning agent. Maintains a distribution over $[k]$, say weights $\alpha_1^t, \dots, \alpha_k^t$ over the agents. Proxy of how poorly they've been doing so far.

$$|\mathrm{L}_{\mathcal{D}_i}(h^t) - \widehat{\mathrm{L}_{\mathcal{D}_i}}(h^t)| \le \epsilon'.$$

$$\epsilon' \ge \quad \frac{1}{T}\sum \mathrm{L}_{P^t}(h^t) \ge \underbrace{\max_{i \in [k]} \frac{1}{T}\sum \mathrm{L}_{\mathcal{D}_i}(h^t) - \overbrace{\frac{\sqrt{T.\log(k)}}{T}}^{\le \epsilon' \text{ for } T = \frac{\log(k)}{\epsilon'^2}}}.$$

$\max_{i \in [k]} \mathrm{L}_{\mathcal{D}_i}(\bar{h}_T)$ when $\bar{h}_T$ is a randomized classifier uniformly from $h^1, \dots, h^T$

(Better) Dynamics 3: No-Regret versus No-Regret

Why we used best-response:

→Adaptivity of the sequence meant that performance is correlated across time.

→ERM ensured that every step approximates the no-regret dynamics on the true game (with expected losses).

An alternative: Picking h^t before looking at $(x_t, y_t) \sim P^t$ gives unbiased estimates.

Convergence rate to multi-objective solution $\approx$ Empirical Regret Min-Player $+$ Empirical Regret Max-Player [H, Jordan, Zhao '22]

Full info: $\sqrt{\log(|H|)/T}$ Bandit: $\sqrt{k/T}$

There is an algorithm for multi-distribution learning with sample complexity
$\tilde{O}\left(\epsilon^{-2}(\log(|\mathrm{H}|) + k\ln(k))\right)$. [H, Jordan, Zhao '22]

Important Message

Online Learning as a Powerful Medium
for Interactions in Learning
(beyond adversarial)

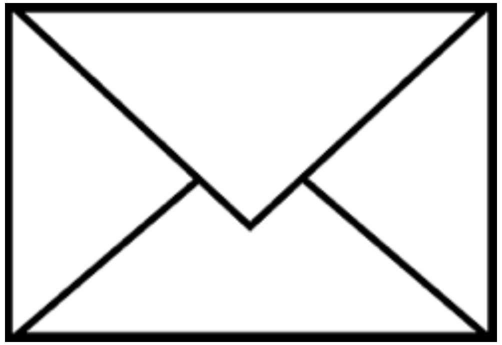

Foundations of
Multi-Objective Statistical Learning

Rest of this talk:

- Multi-Objective Learning, one unifying framework
  1. Collaborative and multi-distribution learning
  2. Multi-calibration
  3. Optimal tradeoffs

Recall Multi-Calibrated predictions

Consider representative subpopulations $\mathcal{S} \subseteq 2^X$. A predictor is multi-calibrated, if for every group $S \in \mathcal{S}$ and every predicted value v

$$\mathbb{E}[y|p(x) \approx v, x \in S] \approx v$$

E.g., if p predicts heart failure probability in a patient, then for each S among patients of S where p predicts 0.1, 10% truly develop heart failure.

Multi-Calibration as Multi-Objective

Consider the class of all predictor $H = [0,1]^X$, the set of loss functions $\ell^{v,S,\sigma}(x, y, p) = \sigma\big(y - p(x)\big)1(p(x) \approx v)\ 1(x \in S)$ for $\sigma \in \{-1, +1\}$, predicted values v and subgroups $S \in \mathcal{S}$, and a single distribution $\mathcal{D}$.

p is $(\mathcal{S}, \epsilon)$-multicalibrated if

$$\max_{\sigma,v,S} \ \mathrm{L}_{\mathcal{D}}^{\sigma,v,S}(p) \ - \overbrace{\min_{p^* \in H} \max_{\sigma,v,S} \mathrm{L}_{\mathcal{D}}^{\sigma,v,S}(p^*)}^{= 0 \text{ by Bayes classifier}} \leq \epsilon$$

Multi-Calibration and learning Challenges

Implement no-regret dynamics, while only sample from $\mathcal{D}$ (can't see $L_{\mathcal{D}}^{v_t,S_t}(p_t)$ exactly).

- For the minimizing player (player learning p_t), can we

→ (best response): Compute p_t with calibration error on (a random) (v_t, S_t) is near 0.

→ Do we need many samples from $\mathcal{D}$?

→(No-regret): Compute p_t with no-regret calibration error against adaptive adversarial choice of $(v_1, S_1), \ldots, (v_T, S_T)$.

→How many samples from $\mathcal{D}$? Can the choice of p_t be deterministic?

→How large is the regret, given that number of predictor $\mathrm{H} = [0,1]^X$ is large.

→For the maximizing player, can we play no-regret with small # samples from $\mathcal{D}$.

- Ideally, we want a single predictor p to be multi-calibrated, instead of the time-averaged distribution over predictors $p_1, \ldots, p_T$. Also, want easier-to-interpret ps.

Different Dynamics and their Tradeoffs

1. For higher moments, having a single choice p matters.
e.g., heart-failure risk being 1 standard deviations from mean.

2. Interpretable/succinct predictors: $p(x)$ is interpreted using memberships of $S \ni x$.
Fewer demographic memberships need to be accessed and recalled to describe p.

3. Evolution of subpopulations: We might not know the populations S and important predicted values v at first but pick up on them as the need arises.

No-regret v No-regret	**No-Regret v Best Response**	**Best Response v No-Regret**
small # samples	medium # samples	small # samples
Distribution over predictors	Single predictor	Distribution over predictors
Interpretable predictor*	More complex predictors	More complex predictors
Not robust	Not robust	Robust to evolution of subpopulations

Multi-calibrated Best Response and No-Regret

Best response: For any distribution over $\ell^{(v,S)}$, there is (randomized) predictor p independent of $\mathcal{D}$ that minimizes the $\ell^{(v,S)}$ calibration loss.

[e.g., intuition by Hart's minmax, or Foster & Vohra '97, **H**, Jordan, Zhao '23]

No-Regret: For any adversarial sequence of losses $\ell^{(v,S)}$, there is an explicit construction for p_t that is 1) deterministic, 2) takes no samples from $\mathcal{D}$, and 3) has regret $\sqrt{T \ln(|\#labels|)}$.

[**H**, Jordan, Zhao '23]

Important bits to remember from this:

- No-regret and Best-response both can be used without samples.
- No-regret uses deterministic choice of a predictor, best response uses a distribution over predictors.
- Regret is small: convergence can be very fast.

Dynamics 1: Best-Response vs No-Regret

Maximizing player choosing a population S_t predicted value v_t and penalizing over/under prediction σ_t uses any no-regret algorithm with regret $\sqrt{T \ln(\#action)} \approx \sqrt{T \ln(|S|/\lambda)}$.

Granularity of the prediction bucket

The minimizing player predicts using a (randomized) best-response p_t for that subpopulation.

The maximizing player estimates $L_{\mathcal{D}}^{\sigma_t, v_t, S_t}(p_t)$ using a single sample $(x, y) \sim \mathcal{D}$.

Then the randomized predictor $\bar{p} \sim Unif(p_1, \dots, p_T)$ is $(\mathcal{S}, \epsilon)$-multi-calibrated after $T = \ln\left(\frac{S}{\lambda}\right)/\epsilon^2$.

Tradeoffs: Non-deterministic predictor, very fast convergence.

Dynamics 2: No-Regret vs Best-Response

Maximizing player choosing a poorly calibrated population S_t predicted value v_t and penalizing over/under prediction σ_t uses best-response.

The minimizing player predicts using a deterministic no-regret p_t without sampling.

Then the randomized predictor $\bar{p} \sim Unif(p_1, \ldots, p_T)$ is $(\mathcal{S}, \epsilon)$-multi-calibrated after $T = \ln(k)/\epsilon^2$ rounds. In fact, because p_ts are deterministic, at least one of them is also *a deterministic* $(\mathcal{S}, \epsilon)$-multi-calibrated!

Tradeoffs: Deterministic predictor, very fast convergence.

Important Message

Dynamics in Multi-Objective Learning unifies approaches to fairness, collaboration, and robustness.

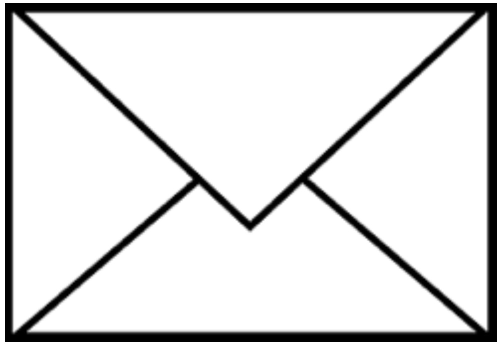

Foundations of Multi-Objective Statistical Learning

Rest of this talk:

- Multi-Objective Learning, one unifying framework
  1. Collaborative and multi-distribution learning
  2. Multi-calibration
  3. Performance-Cost tradeoffs

Beyond Accuracy Guarantees

Agents also incur cost for collecting information:

- E.g., cost for data set curation, privacy cost, etc.
- The protocol shouldn't ask for "unreasonable" amount of data.

→Data-sharing and collaborative learning should be beneficial to all of its users.

How should we procure resources
needed for learning?

Achieve desirable
per-agent tradeoff between accuracy
and sample complexity

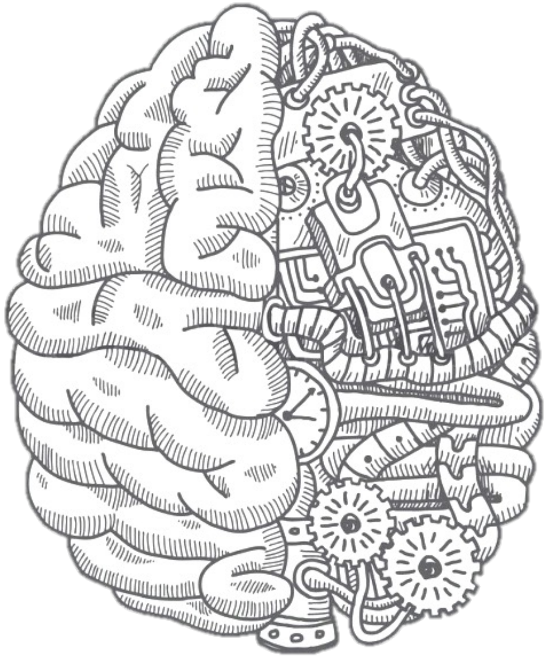

Reasonable Share of Data

<u>**Unreasonable**</u> for agent i if

- We ask i for more data than necessary, if he were to learn by himself.

  → Call avoiding this as **individually rationality.**

- Part of i's contribution is exclusively used to meet the accuracy constraint of other agents and did not affect agent i.

  → Call avoiding this as **stability / an equilibrium.**

[Blum, **H**, Phillips, Shao '21]

State-of-the-art algorithms have poor tradeoffs

Per-agent tradeoff between prediction quality and information (sample complexity):

- Receiving a reasonable return in what resources you put in.

Usability and stability of systems over time:

- Even a small reduction in contribution across the agents impacts algorithmic performance.

60% of agents can unilaterally reduce their contributions to 5% of current levels.

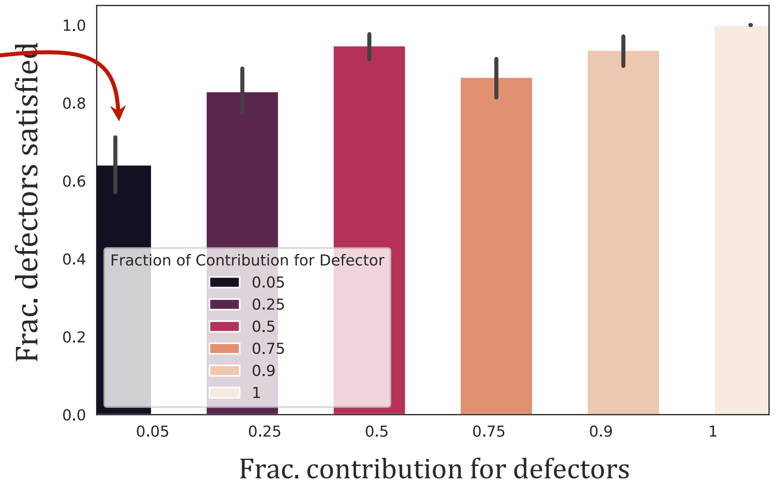

First steps towards using economic theory of incentives (individual rationality, equilibria) to formulate desirable tradeoffs [Blum, **H**, Phillips, Shao '21]

How do Dynamics Impact these tradeoffs?

We don’t know the answers yet formally.

Observationally:

- Non-interactive solutions have much worst tradeoffs than interactive ones.
- Among interactive ones:
    - There is some reason to believe that No-regret v No-Regret has better tradeoffs.
    - Partly because it takes few samples at any time but iterate more over agents.

Are Rationality and Equilibria Restrictive?

Individually rational allocation always exists.

Stable allocations exists, under some natural assumptions.

May require more resources than the optimal collaboration!
→ Rational or stable allocations can be very far from optimal.

[Blum, **H**, Phillips, Shao ’21]

Price of Rationality and Stability

Individually rational or stable equilibria, require more collective resources than the optimal collaboration.

Optimal:  does all the work, others do nothing.

Stable/Rational:  does (almost) no work. Other agents have to do the work.

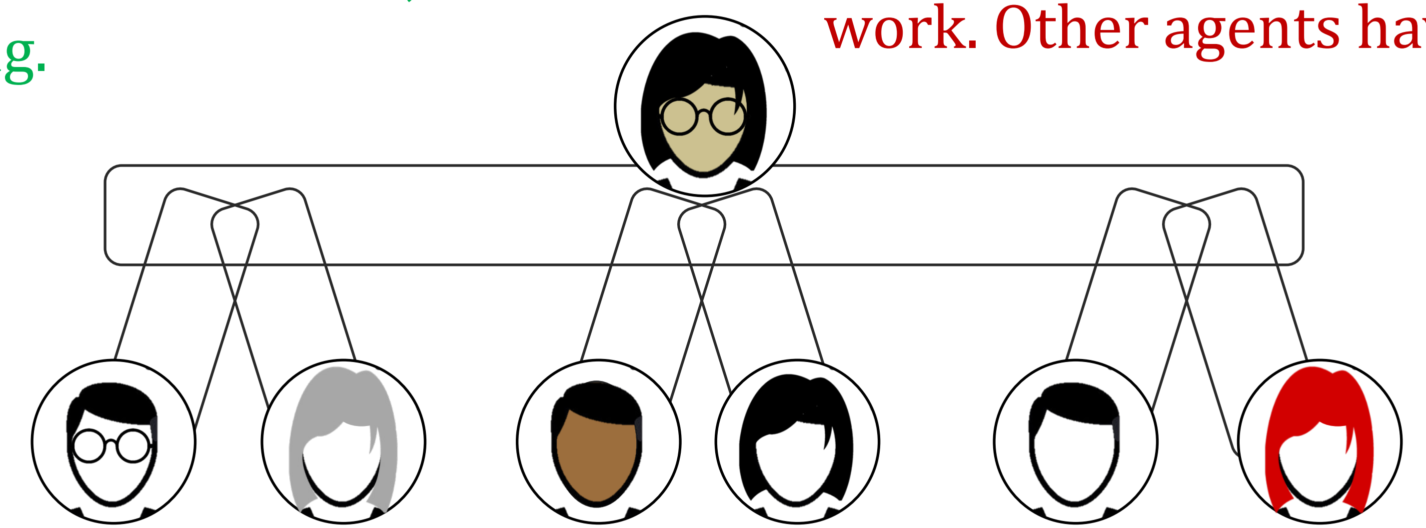

Equilibrium/Individual Rationality: Total work required to be done by other agents is large.

Overall # samples in the best IR/Stable allocation $= \Omega(\sqrt{k}) \times$ Overall # samples in the optimal collaboration

Optimality, Equilibria, and Free Riding

Stable collaborations:

Judiciously introduce small inefficiencies, so everyone can continue benefitting from the system.

Free-riding is a form of necessary inefficiency.

→ Some agents must contribute 0 samples in any stable collaboration.

Is free-riding all that bad?

→ Free-riding must be part of the system, but it's not bad!

→ Apart from free-riders, other agents collaborate optimally.

→ Free-riders don't fundamentally change the optimal collaboration structure between participating agents.

Important Takeaway

New mathematical foundation for

multi-agent statistical and computational learning

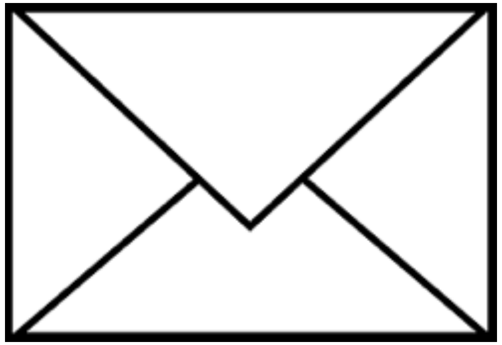

Important Takeaway

New mathematical foundation for

multi-agent statistical and computational learning

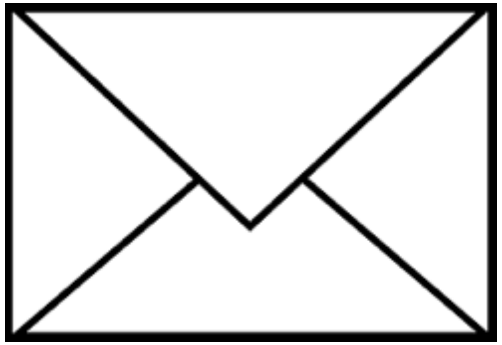

Game theory and online decision making are powerful tools for

for considering per-agent incentives, objectives, and tradeoffs.