Uncertainty and the Social Planner's Problem: Why Sample Complexity Matters

Cyrus Cousins

Brown University
Department of Computer Science
Providence, RI

cyrus_cousins@brown.edu
http://cs.brown.edu/people/ccousins/projects/fairness/home.html

FAccT 서울 2022

Sampling Models for Group-Centric Fair Learning

- Group-centric fair learning considers the *input* and *perspective* of *multiple groups*
  - WLOG assume a set $\mathcal{Z}$ of g groups, i.e., $z \in 1, \ldots, g$
  - Want to learn a mapping $h \in \mathcal{H} \subseteq \mathcal{X} \to \mathcal{Y}$, i.e., from domain $\mathcal{X}$ onto codomain $\mathcal{Y}$
  - Supervised learning process observes $(\mathcal{X}, \mathcal{Y})$ pairs *for each group* $z \in \mathcal{Z}$
- Sampling with multiple groups raises many questions:
  - How is data collected? ♠ What is the cost? ♠ How to measure sample complexity?
- We introduce three models of sampling, and discuss learning in each:
  1. *Joint Sampling*: Each i.i.d. sample contains information for each group. For example, per-group representatives could be shown a shared $x \in \mathcal{X}$ and asked for their feedback, which would then be used to establish some $\mathcal{Y}_i$ for each group i.
  2. *Mixture Sampling*: For each sample, the data are only relevant to one group, i.e., we randomly sample from a *mixture distribution* over groups.
  3. *Conditional Sampling*: Here we *actively choose* from which group to sample. Natural in *active sampling*, *scientific inquiry*, and *stratified sampling* settings, where initial results guide further study.

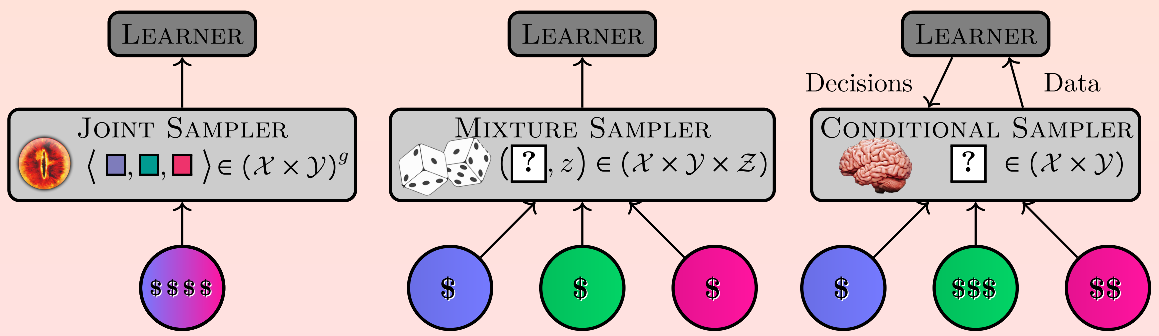

Fair Learning Objectives

- This work generalizes, unifies, and analyzes three disparate fairness concepts
  - Welfare $\mathrm{W}(\cdot; \boldsymbol{w})$ summarizes overall wellbeing (utility $\mathrm{u}(\cdot,\cdot)$) across groups
    - ♥ Generalizes *utility maximization* to multiple groups

$$h^* \leftarrow \operatorname*{argmax}_{h \in \mathcal{H}} \mathrm{W}\left(j \mapsto \mathop{\mathbb{E}}_{(x,y) \sim \mathcal{D}_j} \left[\mathrm{u}(h(x), y) \right]; \boldsymbol{w} \right)$$

  - Malfare $\mathrm{M}(\cdot; \boldsymbol{w})$ summarizes overall illbeing (loss $\ell(\cdot,\cdot)$)
    - ♥ Generalizes *risk minimization* and *minimax fair learning*

$$h^* \leftarrow \operatorname*{argmin}_{h \in \mathcal{H}} \mathrm{M}\left(j \mapsto \mathop{\mathbb{E}}_{(x,y) \sim \mathcal{D}_j} \left[\ell(h(x), y) \right]; \boldsymbol{w} \right)$$

  - Regret measures the utility or loss $\mathrm{s}(\cdot,\cdot)$ lost by compromising on a *shared solution*
    - ♥ Generalizes *multi-group agnostic PAC learning*
    - ♥ Compare *overall solution* h^* to per-group optimal solutions h_j^*

$$h^* \leftarrow \operatorname*{argmin}_{h \in \mathcal{H}} \mathrm{M}\left(j \mapsto \sup_{h_j^* \in \mathcal{H}} \left| \mathop{\mathbb{E}}_{(x,y) \sim \mathcal{D}_j} \left[\mathrm{s}(h(x), y) \right] - \mathop{\mathbb{E}}_{(x,y) \sim \mathcal{D}_j} \left[\mathrm{s}(h_j^*(x), y) \right] \right|; \boldsymbol{w} \right)$$

- Fairness objectives *mathematically encode* the values of a society
  - Different axiomatizations give rise to different objectives
  - There is no "best" or "most fair" objective
  - Various reasonable welfare $\mathrm{W}(\cdot; \boldsymbol{w})$ and malfare $\mathrm{M}(\cdot; \boldsymbol{w})$ functions
    - ♥ Represent different priorities ♥ Make different tradeoffs

Utilitarian, Egalitarian, and the Power Mean Family

- The power-mean for $p \in \mathbb{R}$ summarizes g *values* $\mathcal{S}_{1:g}$ with *weights* $\boldsymbol{w}_{1:g}$ as

$$\mathrm{M}_{p \neq 0}(\mathcal{S}; \boldsymbol{w}) \doteq \sqrt[p]{\sum_{i=1}^{g} \boldsymbol{w}_i \mathcal{S}_i^p}\ , \quad \mathrm{M}_0(\mathcal{S}; \boldsymbol{w}) \doteq \exp\left(\sum_{i=1}^{g} \boldsymbol{w}_i \log(\mathcal{S}_i^p) \right) = \prod_{i=1}^{g} \mathcal{S}_i^{\boldsymbol{w}_i}$$

- Fair welfare requires $p \leq 1$; extremes are interesting special cases
  - $p = 1$ is *weighted sum* over groups (well-studied case)
  - $p = -\infty$ limit is *minimum* over groups (egalitarian or robust maximization)
- Fair malfare (or regret malfare) requires $p \geq 1$
  - $p = \infty$ limit is *maximum* over groups (egalitarian, minimax fair learning)
- Power-means are:
  1. *Axiomatically Justified*
  2. *Interpretable*
     $\mathrm{M}_p(\mathcal{S}; \boldsymbol{w})$ units match $\mathcal{S}_{1:g}$
  3. *Stochastically Stable*
     (for $p \in [-\infty, 0) \cup [1, \infty]$)

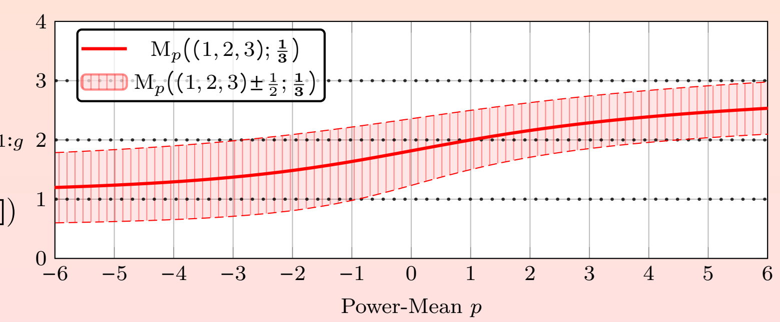

Bounding Generalization Error and Overfitting to Fairness

- Fairness *in trainig* is not sufficient!
  - Less data available for marginalized or minority groups $\Longrightarrow$ *overfitting*
  - Induction bias on welfare, malfare, or regret objectives
- Given (WLOG) some *malfare objective* $\mathrm{M}(\cdot)$, hypothesis space $\mathcal{H}$
  - Exists some *optimal* $h^* \in \mathcal{H}$
  - Want to select (learn) a hypothesis $\hat{h} \in \mathcal{H}$
- $\hat{h}$ should be *almost as good* as h^*
  - ε-δ *Probably Approximately Correct*
  - With probability at least $1 - \delta$ (over training data):

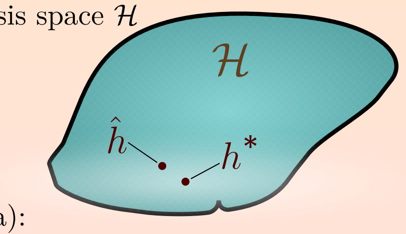

$$\mathrm{M}\left(j \mapsto \mathop{\mathbb{E}}_{(x,y) \sim \mathcal{D}_j} \left[\ell(\hat{h}(x), y) \right]; \boldsymbol{w} \right) \leq \varepsilon + \inf_{h^* \in \mathcal{H}} \mathrm{M}\left(j \mapsto \mathop{\mathbb{E}}_{(x,y) \sim \mathcal{D}_j} \left[\ell(h^*(x), y) \right]; \boldsymbol{w} \right)$$

- Special cases:
  - Utilitarian malfare: **weighted risk minimization**
    - ♥ Minimize *weighted sum* of per-group risks
  - Egalitarian malfare: **minimax fair learning**
    - ♥ Minimize *worst-case* per-group risk

Bernstein-Type Bounds for Malfare Estimation

- Suppose power-mean malfare $\mathrm{M}_p(\cdot; \boldsymbol{w})$ with $p \geq 1$
- Suppose loss range $[0, r]$ and maximum variance $v \doteq \sup_{j \in \mathcal{Z}} \mathop{\mathbb{E}}_{(x,y) \sim \mathcal{D}_j} \left[\ell(h(x), y) \right]$
- Gap between empirical malfare $\hat{\mathrm{M}}$ and true malfare M is bounded as
  1. $\mathbb{P}\left(\left| \mathrm{M} - \hat{\mathrm{M}} \right| \geq \frac{r \ln \frac{2g}{\delta}}{3m} + \sqrt{\frac{2v \ln \frac{2g}{\delta}}{m}} \right) \leq \delta$
  2. $\left| \mathbb{E}[\mathrm{M}] - \mathbb{E}[\hat{\mathrm{M}}] \right| \leq \mathbb{E}\left[\left| \mathrm{M} - \hat{\mathrm{M}} \right| \right] \leq \frac{r \ln(2eg)}{3m} + \sqrt{\frac{2v \ln(2eg)}{m}}$

The Incremental Knowledge Gain of a Single Sample

- Goal is to estimate or optimize the objective to within ε additive error
- How much will an additional sample for group i improve confidence bounds?
- For power-mean malfare, we can cleanly approximate this quantity:
  - Suppose power-mean malfare $\mathrm{M}_p(\cdot; \boldsymbol{w})$ and let $\hat{\mathrm{M}}$ be the *empirical malfare*
  - Let $\hat{\varepsilon}$ denote *confidence interval radius* for group i
  - Let $\hat{\mathrm{M}}^{\uparrow}$ and $\tilde{\mathrm{M}}^{\uparrow}$ be UCB estimates of M with with sample sizes $\boldsymbol{m}_{1:g}$ and $\boldsymbol{m} + \mathbb{1}_i$
  - Then the *incremental impact* of sampling from group i is approximately

$$\hat{\mathrm{M}}^{\uparrow} - \tilde{\mathrm{M}}^{\uparrow} \approx \frac{\hat{\varepsilon}_i \boldsymbol{w}_i}{2\boldsymbol{m}_i + \frac{3}{2}} \left(\frac{\hat{\mathbb{E}}_{\boldsymbol{x}_{i,:}, \boldsymbol{y}_{i,:}}[\ell \circ \hat{h}] + \hat{\varepsilon}_i}{\hat{\mathrm{M}}^{\uparrow}} \right)^{p-1} \approx \frac{\hat{\varepsilon}_i \boldsymbol{w}_i}{2\boldsymbol{m}_i} \left(\frac{\hat{\mathbb{E}}_{\boldsymbol{x}_{i,:}, \boldsymbol{y}_{i,:}}[\ell \circ \hat{h}]}{\hat{\mathrm{M}}} \right)^{p-1}$$

  1. *Inversely proportional* to the amount of effort $\boldsymbol{m}_i$ already spent studying group i
  2. *Proportional* to the current bound radius $\hat{\varepsilon}_i$ and the group weight $\boldsymbol{w}_i$
  3. *Proportional* to the ratio between group risk and $\hat{\mathrm{M}}$ (relative risk)
     (a) Raising this term to the $(p-1)$th power nonlinearly adjusts its impact
     (b) Higher p saturate high-risk groups, tending towards *egalitarianism*
     (c) Decreasing $p \to 1$ takes this term to 1 (constant), tending toward *utilitarianism*

Example: Optimal Sampling under Parametric Gaussian Assumption

- Suppose Gaussian uncertainty over group 1 and group 2 risk values

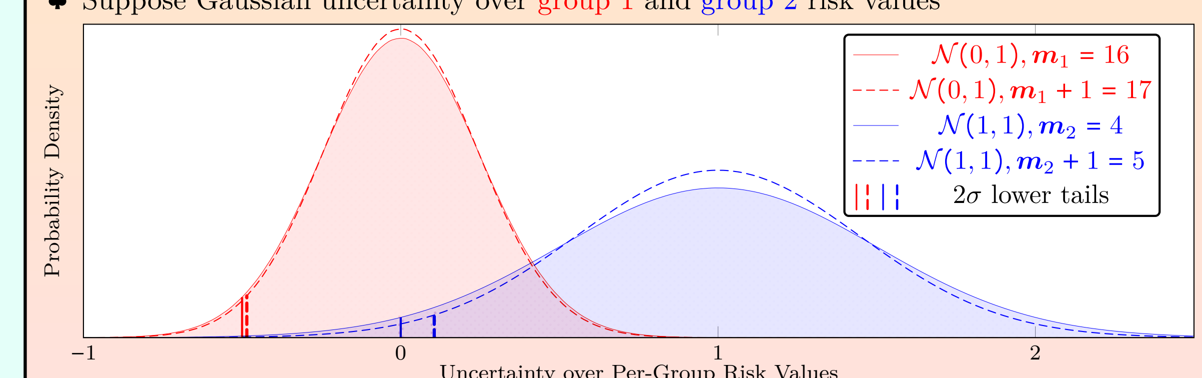

- Optimal choice depends on both *per-group uncertainty* and *objective*
  - Egilitarian malfare: sample group 1, more likely to be the minimum
  - Utilitarian malfare: sample group 2, expect more improvement

Progressive and Active Sampling Algorithms for Fair Learning

- *Progressive sampling* turns *statistical bounds* into *approximation algorithms*
- The basic idea is quite simple:
  1. Start with a small sample from each group
  2. Optimize or estimate the objective on the current sample
  3. Terminate if some optimality condition is met
  4. Draw a larger sample and repeat from (2)
- We can estimate any continuous monotonic fairness objective
  - No continuity $\Longrightarrow$ algorithm may never terminate
  - Continuity $\Longrightarrow$ eventual termination under *infinite sampling schedule*
  - Lipschitz continuity $\Longrightarrow$ sufficient *finite sampling schedule* (more efficient)
- Efficiently operate under various sampling models
  - Joint Sampling, Mixture Sampling: only decision is when to terminate
  - Conditional Sampling: must also decide where to sample!
    - ♥ Active learning with *greedy optimality heuristic*:
      - ♦ Balance *cost* and *estimated bound improvement*