Announcement

- Grades for HW2 and project proposal are released

CS6501: Topics in Learning and Game Theory
(Fall 2019)

Learning from Strategically Transformed Samples

Instructor: Haifeng Xu

Part of the Slides are provided by Hanrui Zhang

Outline

- Introduction
- The Model and Results

Signaling

Q: Why attending good universities?

Q: Why publishing and presenting at top conferences?

Q: Why doing internships?

Signaling

Q: Why attending good universities?

Q: Why publishing and presenting at top conferences?

Q: Why doing internships?

- All in all, these are just signals (directly observable) to indicate "excellence" (not directly observable)

Signaling

Q: Why attending good universities?

Q: Why publishing and presenting at top conferences?

Q: Why doing internships?

- All in all, these are just signals (directly observable) to indicate "excellence" (not directly observable)
- Asymmetric information between employees and employers

JOB MARKET SIGNALING *

MICHAEL SPENCE

2001 Nobel Econ Price is awarded to research on asymmetric information

Signaling

- A simple example
  - We want to hire an Applied ML researcher
  - Only two types of ML researchers in this world
  - Easy to tell

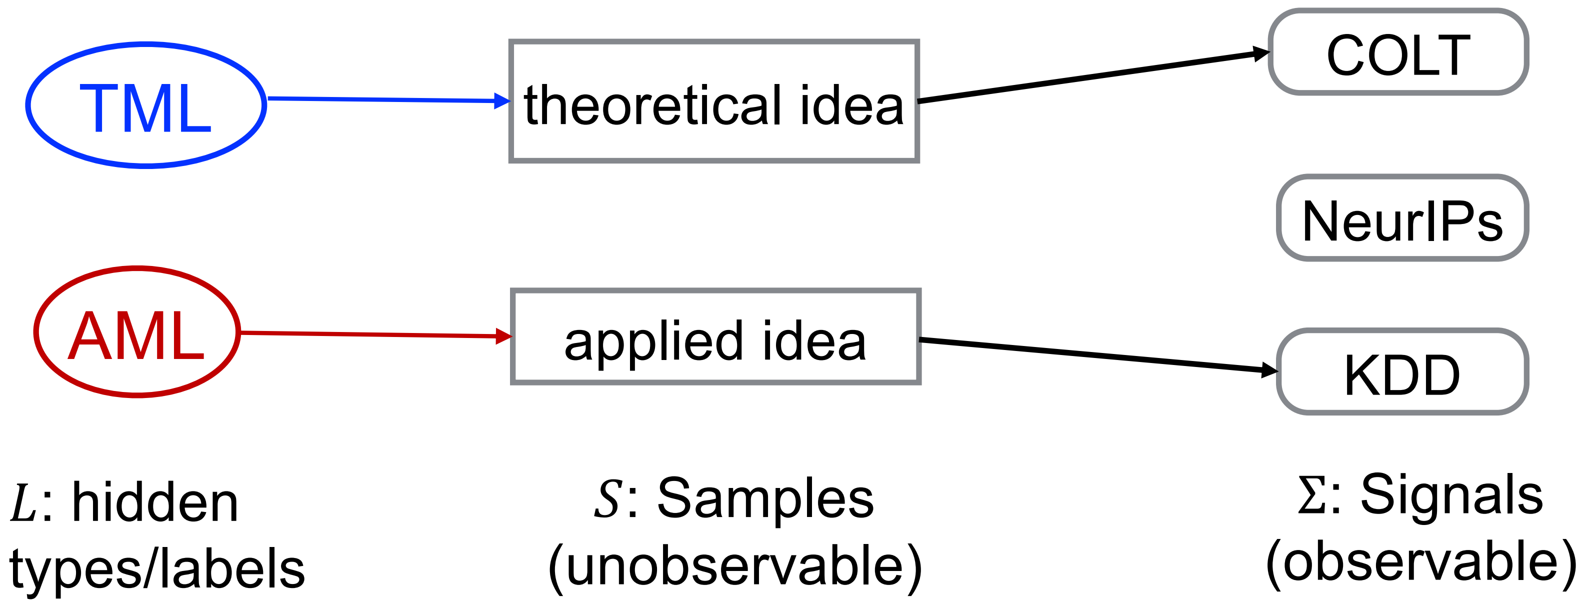

Signaling

- A simple example
  - We want to hire an Applied ML researcher
  - Only two types of ML researchers in this world
  - Easy to tell

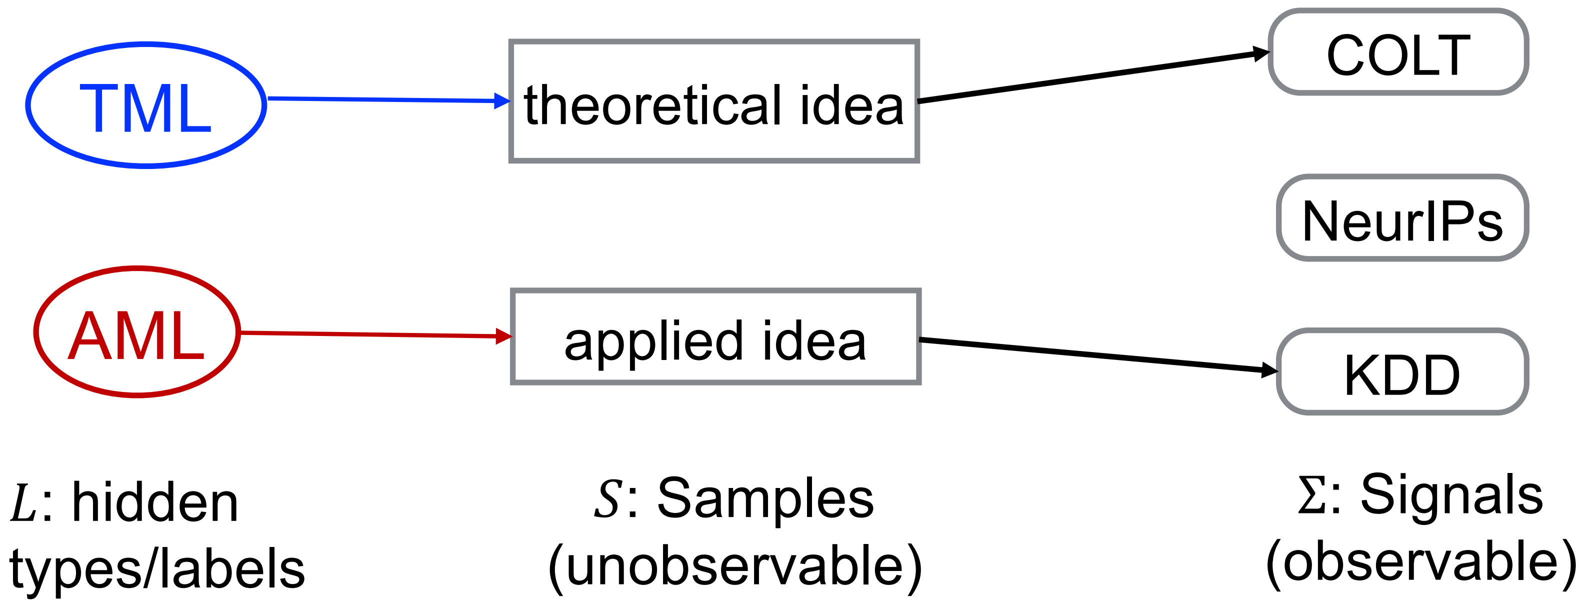

Our world is known to be noisy….

Signaling

➢ A simple example
- We want to hire an Applied ML researcher
- Only two types of ML researchers in this world

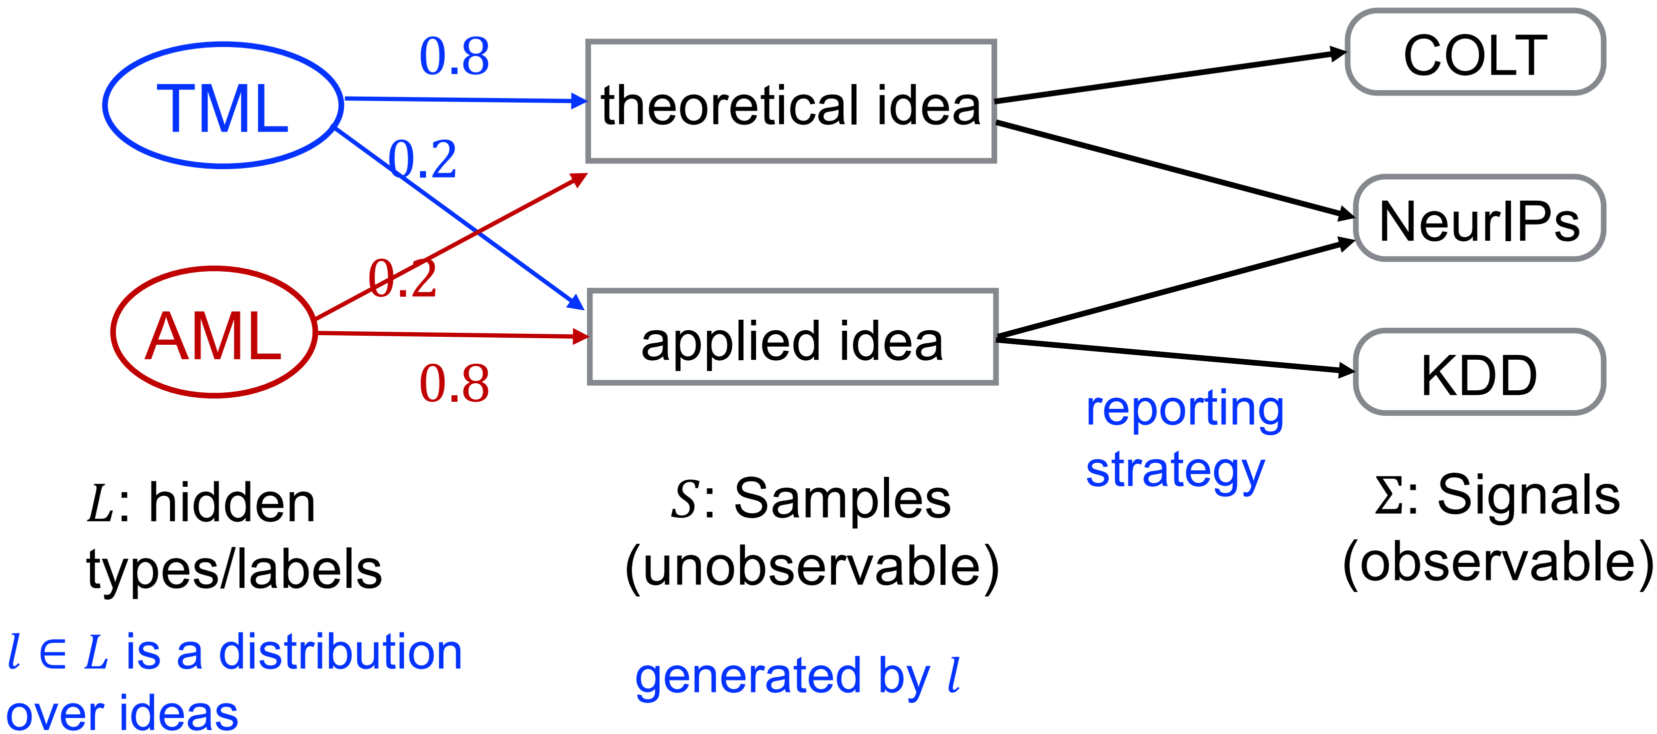

Signaling

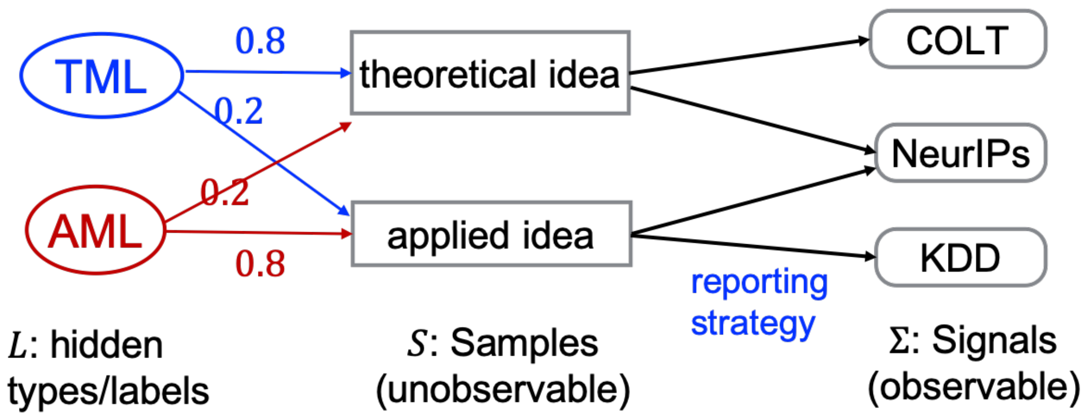

- Agent's problem:
  - How do I distinguish myself from other types?
  - How many ideas do I need for that?
- Principle's problem:
  - How do I tell AML agents from others (a classification problem)?
  - How many papers should I expect to read?

Answers for this particular instance?

Signaling

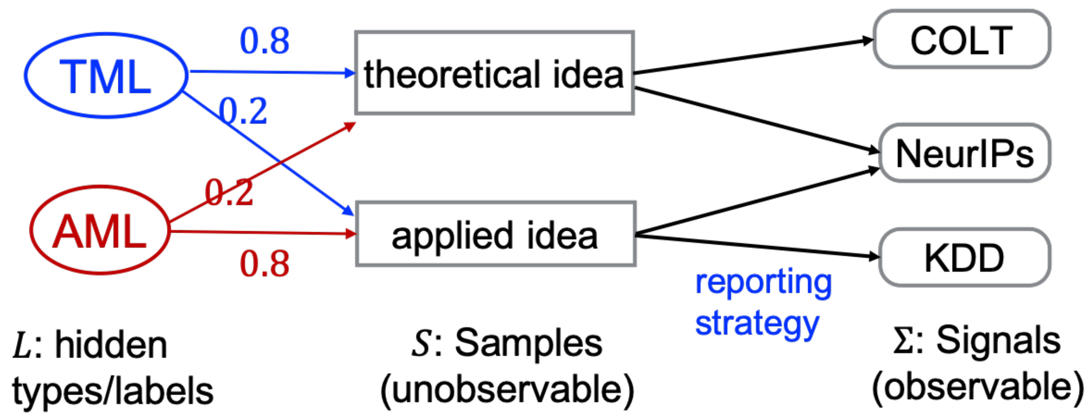

- Agent's problem:
  - How do I distinguish myself from other types?
  - How many ideas do I need for that?
- Principle's problem:
  - How do I tell AML agents from others (a classification problem)?
  - How many papers should I expect to read?

Generally, classification with strategically transformed samples

What Instances May Be Difficult?

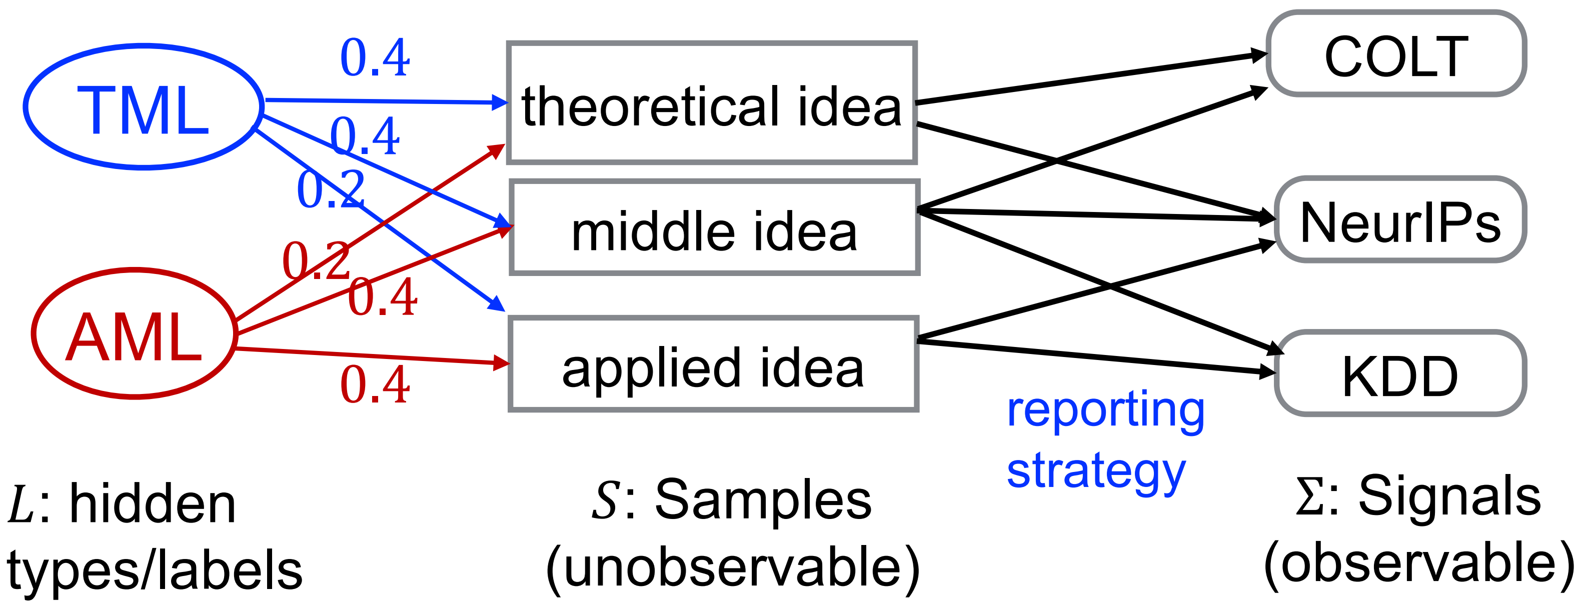

Intuitions

- Agent: try to report as far from others as possible
- Principal: examine a set of signals that maximally separate AML from TML

Outline

- Introduction
- The Model and Results

Model

- Two distribution types/labels: $l \in \{g, b\}$
  - g should be interpreted as "desired", not necessarily good or bad
- $g, b \in \Delta(S)$ where S is the set of samples
- Bipartite graph $G = (S \cup \Sigma, E)$ captures feasible signals for each sample: $(s, \sigma) \in E$ iff σ is a valid signal for s
- g, b, G publicly known; S, Σ both discrete
- Distribution $l \in \{g, b\}$ generates T samples

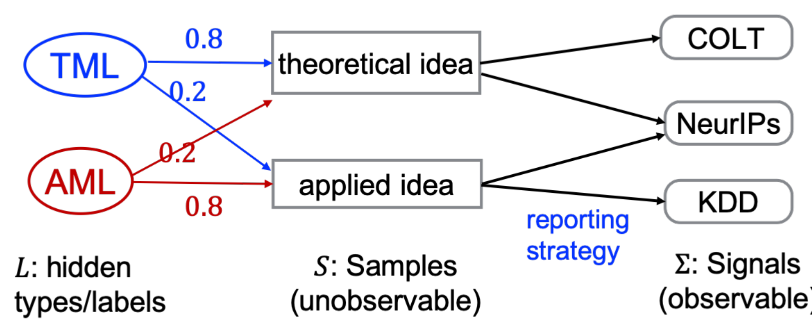

Model

- Two distribution types/labels: $l \in \{g, b\}$
  - g should be interpreted as "desired", not necessarily good or bad
- $g, b \in \Delta(S)$ where S is the set of samples
- Bipartite graph $G = (S \cup \Sigma, E)$ captures feasible signals for each sample: $(s, \sigma) \in E$ iff σ is a valid signal for s
- g, b, G publicly known; S, Σ both discrete
- Distribution $l \in \{g, b\}$ generates T samples
- A few special cases
  - Agent can hide samples, as in last lecture (captured by adding a "empty signal")
  - Signal space may be the same as samples (i.e., $S = \Sigma$); G captures feasible "lies"

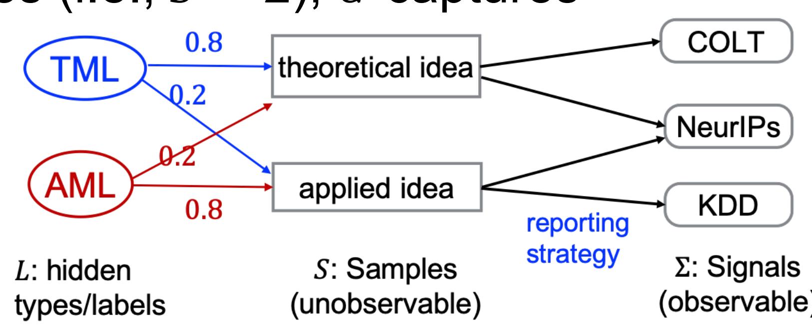

The Game

Agent's reporting strategy π transform T samples to a set R of T signals

- A reporting strategy is a signaling scheme
  - Fully described by $\pi(\sigma|s)$ = prob of sending signal σ for sample s
  - $\sum_\sigma \pi(\sigma|s) = 1$ for all s

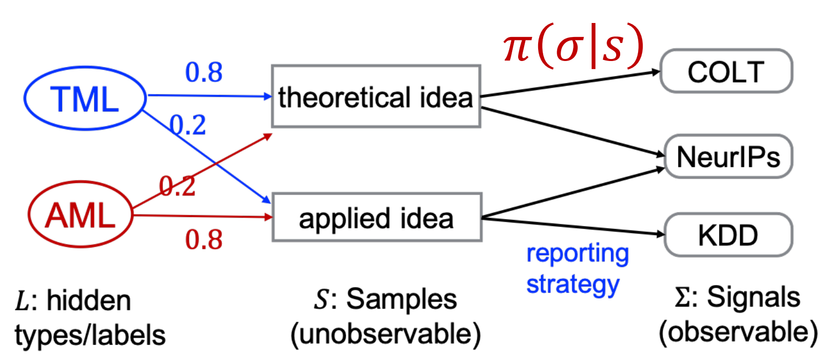

The Game

Agent's reporting strategy π transform T samples to a set R of T signals

- A reporting strategy is a signaling scheme
  - Fully described by $\pi(\sigma|s)$ = prob of sending signal σ for sample s
  - $\sum_\sigma \pi(\sigma|s) = 1$ for all s
- Given T samples, π generates T signals (possibly randomly) as an agent report $R \in \Sigma^T$
- A special case is deterministic reporting strategy

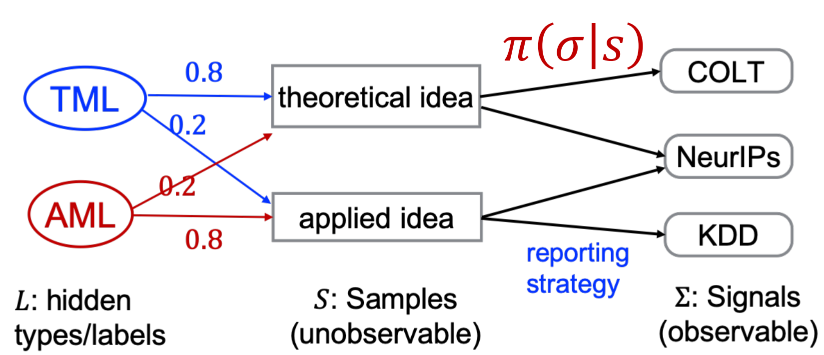

The Game

Agent's reporting strategy π transform T samples to a set R of T signals

- Objective: maximize probability of being accepted

Principal's action $f: \Sigma^T \to [0,1]$ maps agent's report to an acceptance prob

- Objective: minimize prob of mistakes (i.e., reject g or accept b)

Remark:

- Timeline: principal announces f first; agent then best responds
- Type g's [b's] incentive is aligned with [opposite to] principal

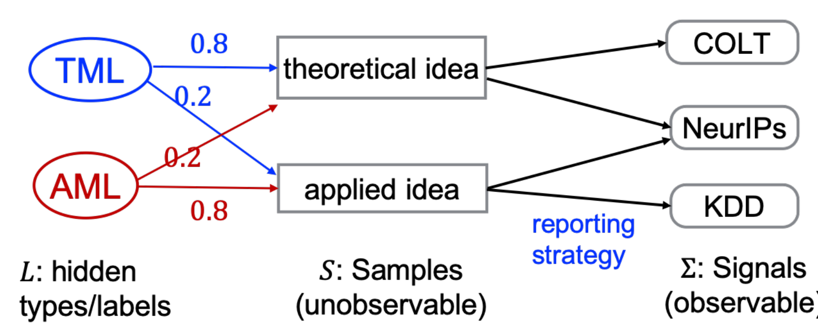

A Simpler Case

- Say $l \in \{g, b\}$ generates $T = \infty$ many samples
- Any reporting strategy π generates a distribution over Σ
  - $\Pr(\sigma) = \sum_{s \in S} \pi(\sigma|s) \cdot l(s) = \pi(\sigma|l)$ (slight abuse of notation)
  - $\pi(\sigma|l)$ is linear in variables $\pi(\sigma|s)$
- Intuitively, type g should make his π "far from" other's distribution
  - Total variance (TV) distance turns out to be the right measure

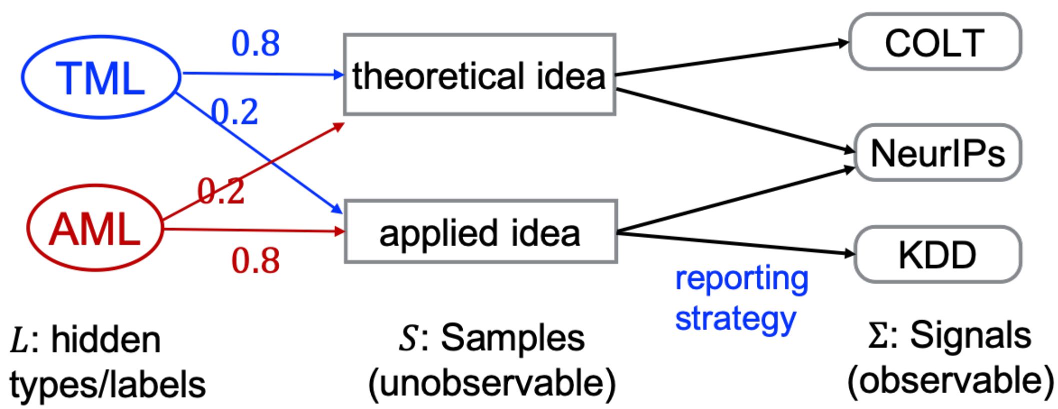

Total Variance Distance

- Discrete distribution x, y supported on Σ
  - Let $x(A) = \sum_{\sigma \in A} x(\sigma) = \Pr_{\sigma \sim x}(\sigma \in A)$

$$d_{TV}(x,y) = \max_{A}[x(A) - y(A)]$$

$$= \sum_{\sigma:\, x(\sigma)>y(\sigma)}[x(\sigma) - y(\sigma)]$$

$$= \frac{1}{2}\underbrace{\sum_{\sigma:\, x(\sigma)>y(\sigma)}[x(\sigma) - y(\sigma)]} + \frac{1}{2}\underbrace{\sum_{\sigma: y(\sigma)\geq x(\sigma)}[y(\sigma) - x(\sigma)]}$$

These two terms are equal

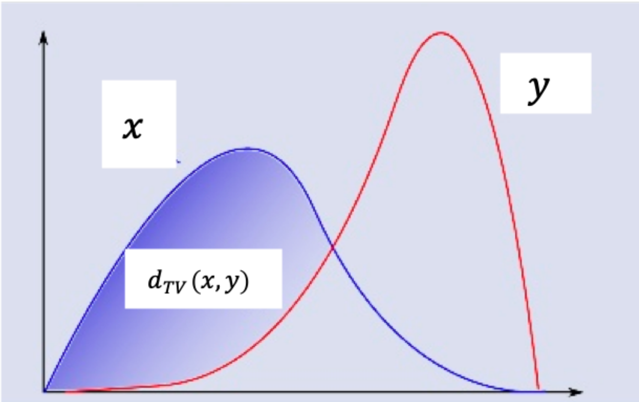

Total Variance Distance

- Discrete distribution x, y supported on Σ
  - Let $x(A) = \sum_{\sigma \in A} x(\sigma) = \Pr_{\sigma \sim x}(\sigma \in A)$

$$d_{TV}(x,y) = \max_{A}[x(A) - y(A)]$$

$$= \sum_{\sigma:\, x(\sigma) > y(\sigma)} [x(\sigma) - y(\sigma)]$$

$$= \frac{1}{2}\sum_{\sigma:\, x(\sigma) > y(\sigma)} [x(\sigma) - y(\sigma)] + \frac{1}{2}\sum_{\sigma: y(\sigma) \geq x(\sigma)} [y(\sigma) - x(\sigma)]$$

$$= \frac{1}{2}\sum_{\sigma} |x(\sigma) - y(\sigma)|$$

$$= \frac{1}{2}||x - y||_1$$

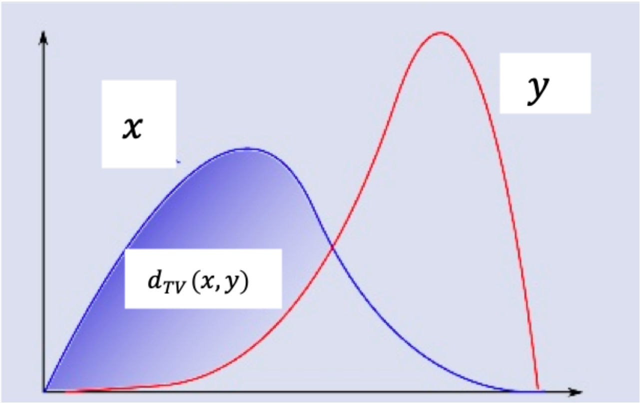

How Can g Distinguish Himself from b?

- Type g uses reporting strategy π (and b uses ϕ)
- Type g wants $\pi(\cdot|g)$ to be far from $\phi(\cdot|b)$ → What about type b?
- This naturally motivates a zero-sum game between g, b

$$\max_{\pi} \min_{\phi} d_{TV}\left(\pi(\cdot|g), \phi(\cdot|b)\right) = d_{DTV}(g, b)$$

Game value of this zero-sum game

How Can g Distinguish Himself from b?

- Type g uses reporting strategy π (and b uses ϕ)
- Type g wants $\pi(\cdot|g)$ to be far from $\phi(\cdot|b)$ → What about type b?
- This naturally motivates a zero-sum game between g, b

$$\max_{\pi} \min_{\phi} d_{TV}\left(\pi(\cdot|g), \phi(\cdot|b)\right) = d_{DTV}(g, b)$$

Note $d_{DTV}(g, b) \geq 0$….now, what happens if $d_{DTV}(g, b) > 0$?

How Can g Distinguish Himself from b?

- Type g uses reporting strategy π (and b uses ϕ)
- Type g wants $\pi(\cdot\,|g)$ to be far from $\phi(\cdot\,|b)$ → What about type b?
- This naturally motivates a zero-sum game between g, b

$$\max_{\pi} \min_{\phi} d_{TV}\left(\pi(\cdot\,|g), \phi(\cdot\,|b) \right) = d_{DTV}(g, b)$$

Note $d_{DTV}(g, b) \geq 0$….now, what happens if $d_{DTV}(g, b) > 0$?

- g has a strategy π^* such that $\mathrm{d_{TV}}\left(\pi^*(\cdot\,|g), \phi(\cdot\,|b)\right) > 0$ for any ϕ
- Using π^*, g can distinguish himself from b with constant probability via $\Theta\left(\frac{1}{\left(d_{DTV}(g,b)\right)^2}\right)$ samples
  - Recall: $\Theta(\frac{1}{\epsilon^2})$ samples suffice to distinguish x, y with $d_{TV}(x, y) = \epsilon$
  - Principal only needs to check whether report R is drawn from $\pi^*(\cdot\,|g)$ or not

How Can g Distinguish Himself from b?

- So $d_{DTV}(g, b) > 0$ is sufficient for distinguishing g from b
- It turns out that it is also necessary

Theorem:

1. If $d_{DTV}(g, b) = \epsilon > 0$, then there is a policy f that makes mistakes with probability δ when #samples $T \geq 2 \ln\left(\frac{1}{\delta}\right) / \epsilon^2$.
2. If $d_{DTV}(g, b) = 0$, then no policy f can separate g from b regardless how large is #samples T.

How Can g Distinguish Himself from b?

- So $d_{DTV}(g, b) > 0$ is sufficient for distinguishing g from b
- It turns out that it is also necessary

Theorem:

1. If $d_{DTV}(g, b) = \epsilon > 0$, then there is a policy f that makes mistakes with probability δ when #samples $T \geq 2 \ln\left(\frac{1}{\delta}\right) / \epsilon^2$.
2. If $d_{DTV}(g, b) = 0$, then no policy f can separate g from b regardless how large is #samples T.

Remarks:

- Prob of mistake δ can be made arbitrarily small with more samples
- We have shown the first part
- Second part is more difficult to prove, uses an elegant result for matching theory

But…Deciding Whether $d_{DTV}(g, b) > 0$ is Hard

Theorem: it is NP-hard to check whether $d_{DTV}(g, b) = 0$ or not.

- Recall $d_{DTV}(g, b) = \max_{\pi} \min_{\phi} d_{TV}\left(\pi(\cdot \mid g), \phi(\cdot \mid b)\right)$

But…Deciding Whether $d_{DTV}(g, b) > 0$ is Hard

Theorem: it is NP-hard to check whether $d_{DTV}(g, b) = 0$ or not.

- Recall $d_{DTV}(g, b) = \max_{\pi} \min_{\phi} d_{TV}\,(\,\pi(\cdot\,|g)\,, \phi(\cdot\,|b)\,)$
- Wait…this is a zero-sum game, and we can solve it in poly time?

Q: What goes wrong?

But…Deciding Whether $d_{DTV}(g, b) > 0$ is Hard

Theorem: it is NP-hard to check whether $d_{DTV}(g, b) = 0$ or not.

- Recall $d_{DTV}(g, b) = \max_{\pi} \min_{\phi} d_{TV}\left(\pi(\cdot\,|g), \phi(\cdot\,|b)\right)$
- Wait…this is a zero-sum game, and we can solve it in poly time?

Q: What goes wrong?

- We can only solve normal-form zero-sum games in poly time
- In that case, utility fnc is linear in both players' strategies
  - Can generalize to concave-convex utility fnc
  - But here, utility fnc is convex in both player's strategies

But…Deciding Whether $d_{DTV}(g, b) > 0$ is Hard

Theorem: it is NP-hard to check whether $d_{DTV}(g, b) = 0$ or not.

- Recall $d_{DTV}(g, b) = \max_{\pi} \min_{\phi} d_{TV}\left(\pi(\cdot\,|g) , \phi(\cdot\,|b) \right)$

Corollary: it is NP-hard to compute g's best strategy π^*.

Proof:

- Will argue if we can compute π^*, then we can check $d_{DTV}(g, b) = 0$ or not
  - Thus computing π^* must be hard (actually "harder" than checking $d_{DTV}(g, b) = 0$)
- If we computed π^*, to compute $d_{DTV}(g, b)$, we only need to solve $\min_{\phi} d_{TV}\left(\pi^*(\cdot\,|g) , \phi(\cdot\,|b) \right)$ which is convex in ϕ
  - Minimize convex fnc can be done efficiently in poly time (well-known)
- First example of reduction in this class

Some Remarks

- Separability is determined by some "distance" between g, b
  - A generalization of TV distance to strategic setting
  - The principal's policy is relatively simple
  - It is more of our own job to distinguish ourselves from others, rather than the employer's
- The model can be generalized to many "good" (g_i) and "bad"(b_j) distributions
  - Principal wants to accept any g_i and reject any b_j
  - Separability is determined by $\min_{i,j} d_{DTV} (g_i, b_j)$
- The agent's reporting strategy can even be adaptive
  - i.e., the π is different for different samples and may depend on past signals
  - Results do not change

Next Lecture will talk about how to utilize strategic manipulations to induce desirable social outcome

Thank You

Haifeng Xu

University of Virginia

hx4ad@virginia.edu