Peer Consistency Mechanisms

Boi Faltings

AI Laboratory, EPFL
boi.faltings@epfl.ch
http://liawww.epfl.ch/

2016

Peer Consistency Mechanisms

- Scoring rule mechanisms require that ground truth becomes known.
- Alternative idea: use reports of peers as ground truth!
  Reward = Pay(A,g) where g = report of a randomly chosen peer.
- Truthtelling becomes an equilibrium: if peers are truthful, truthtelling is the best response.
- Weaker concept, but more broadly applicable:
  - community sensing.
  - product reviews.
  - preferences, opinions, etc.
- Mechanisms are for reporting a value, not a distribution.

Types of Peer Consistency Mechanisms

There is no peer consistency mechanism for arbitrary belief systems!
Mechanisms depend on belief systems:

- Output Agreement: categorical (uncorrelated) distribution of values.
- Peer Prediction: homogeneous agent population with identical and known belief structure.
- Shadowing Mechanisms: weaken need for common posterior.
- Peer Truth Serum: common prior beliefs, but heterogenous belief updates.

Output Agreement Mechanisms

Term coined by von Ahn for Image Labeler:

- ask people to label an image.
- pay a reward if two people give the same label.

Q: When does this incentivize truthfulness/maximum effort?
A: When agents believe that honest peers are most likely to obtain the same value.

Setting

- agent gets observation o of the image.
- agent submits answer $a \in \{x_1, .., x_N\}$.
- center randomly selects reference report b submitted by a peer.
- center pays agent C if $a = b$, 0 otherwise.

Assume peer is truthful $\Rightarrow$ report a that maximizes $Pr(a|o)$

Limitations of Output Agreement

Many possible labels:

Empire State Building
Manhattan
New York City
Skyscraper
City
America

Equilibrium strategy depends on modeling peer's knowledge and beliefs.
Doesn't always encourage most specific answer.

Limitations of Output Agreement

- Suppose agents report the quality of service of Blue Star Airlines, with very high reputation.
- My plane is late and baggage lost.
- Q: Should I report poor service?
- A: no, because most people enjoy good service, so my report will not match the peer!
- Compensation needs to depend on probability of matching!

Peer Prediction Method

- Rather than reward the most likely value...
- ...reward accurate report of posterior distribution!
- Peer prediction method (MRZ 2005):
  - each value for answer $a = x_i$ is associated with a posterior distribution $\hat{Pr}(x|x_i)$
  - use proper scoring rule to score this posterior against peer report.

Reporting poor service...

With peer prediction method:

- Answer= good service: $\hat{P}r(good) = 0.9, \hat{P}r(bad) = 0.1$
- Answer = bad service: $\hat{P}r(good) = 0.8, \hat{P}r(bad) = 0.2$

$\Rightarrow$

- if 80% (or less) of peer reports are "good", "bad" will be the best answer.
- if 90% (or more) of peer reports are "good", "good" will be the best answer.

Numerical example

- Let prior $Pr(good) = 0.85, Pr(bad) = 0.15$ (85% positive reviews)
- Agent a_i observes bad service.
- With output agreement:
  $E[Payoff("good")] = 0.8, E[Payoff("bad")] = 0.2$
  $\Rightarrow$ best to report "*good*"
- Assume quadratic scoring rule $(2p(x) - \sum p(x)^2)$:

$$
\begin{aligned}
E[Payoff("good")) = 0.8(2 \cdot 0.9 - 0.82) + 0.2(2 \cdot 0.1 - 0.82) &= 0.66 \\
E[Payoff("bad")) = 0.8(2 \cdot 0.8 - 0.68) + 0.2(2 \cdot 0.2 - 0.68) &= 0.68
\end{aligned}
$$

$\Rightarrow$ truthful reporting is more profitable, even though it's not the most likely answer!

Improving peer prediction

2 issues:

- general scoring rules generate inefficient payments.
  $\Rightarrow$ generate simpler and more efficient rules using automated mechanism design.
- uninformative equilibria are focal.
  $\Rightarrow$ score against distribution of multiple peer reports.

Automated mechanism design

We need to find payments $Pay(g,g)$ and $Pay(b,b)$ such that:

$$\begin{aligned} Pr(g|g)Pay(g,g) &> Pr(b|g)Pay(b,b) + \epsilon_g \\ Pr(b|b)Pay(b,b) &> Pr(g|b)Pay(g,g) + \epsilon_b \end{aligned}$$

In this example, assuming $\epsilon_g = \epsilon_b = 0.1$:

$$\begin{aligned} 0.9Pay(g,g) &> 0.1Pay(b,b) + 0.1 \\ 0.2Pay(b,b) &> 0.8Pay(g,g) + 0.1 \end{aligned}$$

Solution by linear program

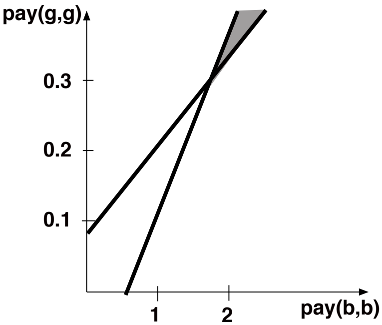

Minimize expected expenditure $Pr(g)Pay(g,g) + Pr(b)Pay(b,b)$,
here: $0.85Pay(g,g) + 0.15Pay(b,b)$
$\Rightarrow$ solution: pay(g,g) $=$ 0.3, pay(b,b)$=$ 1.7
Expected payment $=$ 0.28 (vs. 0.799 w/scoring rule)

Uninformative equilibria

3 pure equilibria:

1. truthful: expected payment $= 0.28$
2. always reporting "good": expected payment $= 0.3$
3. always reporting "bad": expected payment $= 1.7$

$\Rightarrow$ truthfulness is not attractive!

Eliminating uninformative equilibria

Use 3 reference reports and count the number of "g"s:

$pr(\|b=g\|\|o)$	0	1	2	3
b	0.008	0.096	0.384	0.512
g	0.001	0.027	0.243	0.729

$pay(o, \|b=g\|)$	0	1	2	3
b	0	10	0	ϵ
g	ϵ	0	2	0

Truthtelling is a strict equilibrium:

$$\begin{aligned} o = bad &: E[Pay(\text{"}bad\text{"})] = 0.96 &> \quad E[Pay(\text{"}good\text{"})] = 0.768 \\ o = good &: E[Pay(\text{"}bad\text{"})] = 0.27 &< \quad E[Pay(\text{"}good\text{"})] = 0.468 \end{aligned}$$

but all "good" or all "bad" is not a strict or weak equilibrium.

Shadowing Mechanisms

- Peer prediction requires a $\hat{P}r(\cdot|x_i)$ for every value of x_i
- Construct from prior distribution $\hat{P}r(\cdot)$ by letting $\hat{P}r(y|x_i) = \hat{P}r(y) + \delta$ and renormalizing distribution.
- As long as agent's posterior is shifted in the same way, scoring rule will give highest expected reward for a truthful report.

Common Prior Mechanisms

- Peer Prediction requires agents to have common posterior beliefs $\hat{Pr}$ for each measurement.
- They may have very different proficiency and confidence in their observations, making posteriors different.
- However, agents have the same prior information $\Rightarrow$ prior is likely to be the same.
- Can we use the shadowing idea to get a more general mechanism?

Desiderata

- Agents have a common prior Pr.
- Shadow posteriors:

$$\begin{aligned} \hat{P}r(x_i|x_i) &= Pr(x_i) + \delta \\ \hat{P}r(x_j|x_i) &= Pr(x_j) - \delta/(N-1); x_j \neq x_i \end{aligned}$$

- Logarithmic scoring rule $Pay(A, g) = \ln A(g) \Rightarrow$

$$Pay(\hat{P}r(x_i|x_i), x_i) - Pay(Pr(x_i), x_i) \simeq \frac{\delta}{Pr(x_i)}$$

- Reporting randomly according to the prior should have reward 0 $\Rightarrow$ matching peer at value x_i should have reward proportional to $\frac{1}{Pr(x_i)}$.

Peer Truth Serum

- Center knows distribution R; assume agent prior is close to R.
- Reward agreement with peer report on value x_i with $Pay(x_i, x_i) = 1/R(x_i)$, 0 otherwise.
- Incentive Compatibility Condition:

$$\begin{aligned} & E_{Pr(r|x_i)}[Pay(x_i, r)] = Pr(x_i|x_i)Pay(x_i, x_i) = Pr(x_i|x_i)/R(x_i) \\ > \quad & E_{Pr(r|x_i)}[Pay(x_j, r)] = Pr(x_j|x_i)Pay(x_j, x_j) = Pr(x_j|x_i)/R(x_j) \end{aligned}$$

- when $R = Pr$, translates to *self-predicting* condition:

$$\frac{Pr(x_i|x_i)}{Pr(x_i)} > \frac{Pr(x_j|x_i)}{Pr(x_j)}$$

Helpful Reporting

What if $R \neq Pr$ (for example, on initializing the mechanism)?
Consider that Pr is more *informed*, i.e. closer to true distribution Q than R.

$\Rightarrow$ agents partition values into:

- under-reported: $R[x] < Pr[x] \Leftrightarrow R[x] < Q[x]$
- over-reported: $R[x] \geq Pr[x] \Leftrightarrow R[x] \geq Q[x]$

Non-truthful strategy: report x instead of y:

- Always profitable if x under-reported and y over-reported.
- Never profitable if x over-reported and y under-reported

Helpful strategy: never report over-reported x for under-reported y.

Asymptotic Accuracy

- Assume center maintains R as an aggregate over reports received over time (for example histogram).
- Asymptotically accurate: R converges to true distribution Q.
- Theorem: Any mechanism that induces helpful reporting is asymptotically accurate. ($\Leftarrow$ never falsely report over-reported value)
- Theorem: Peer Truth Serum admits equilibria in helpful strategies.

Other equilibria...

- All agents report x with smallest $R[x]$.

$\Rightarrow$ equilibrium with highest possible payoff.

- Will lead to uniformative, uniform distribution.
- However, can be detected: distribution of reports varies a lot over time.
- Can also be thwarted if R is not public.

Peer Truth Serum for Crowdsourcing (PTSC)

- Idea: collect R from agents' reports, but keep it private.
- $R =$ histogram of reports from a set of many *similar* tasks.
- Peer report is chosen from reports on the *same* task.
- Agent should believe that $Pr \simeq R$ (in the limit of infinitely many tasks).
- But that for its own task, $Pr(o|o)/R[o]$ is maximized for its own observation o.

Algorithm (PTSC)

1. collect answers to a set of similar tasks $\mathcal{T}$ from crowdworkers.
2. for worker w, calculate $R_w(x) = \frac{num(x)}{\sum_y num(y)}$, where reports by worker w are excluded.
3. for each task t_w carried out by worker w, select a peer worker p that has solve the same task. If they gave the same answer x, reward w with $\alpha(1/R_w(x) - 1)$, otherwise charge α.

Properties (PTSC)

- truthful equilibrium when agents' beliefs satisfy self-predicting condition.
- expected payoff $= 0$ for random answers according to R
- expected payoff < 0 for random answers according to another distribution.
- truthful equilibrium has the highest payoff.

Example (PTSC)

Task	Answers for the task
t_1	b , a , a , c
t_2	$\boxed{b}$, b , b , a
t_3	a , a , b , a
t_4	a , d , a , a
t_5	c , c , a , b
t_6	d , a , d , d
t_7	a , $\boxed{a}$, c , a
t_8	b , b , a , b
t_9	a , a , a , a
t_{10}	b , b , a , b

Answer	a	b	c	d
Count	20	12	4	4
R	0.50	0.30	0.1	0.1
R_w	0.50	0.29	0.105	0.105

t_7:

honest: E[pay(a)] = 4/3-1 = 1/3

strategic: E[pay(d)] = 0

random:

E[pay] = 1/6-0.3+0.7/3-0.1=0

Self-Predicting Assumption

Correct answer		Observed answer a	b	c	d
a	$Count(a)$	15	2	2	1
	$freq(\cdot \mid a)$	**0.75**	0.1	0.1	0.05
b	$Count(b)$	3	9	0	0
	$freq(\cdot \mid b)$	0.25	**0.75**	0	0
c	$Count(c)$	1	1	2	0
	$freq(\cdot \mid c)$	0.25	0.25	**0.5**	0
d	$Count(d)$	1	0	0	3
	$freq(\cdot \mid d)$	0.25	0	0	**0.75**
	$Count$	20	12	4	4
	R	0.5	0.3	0.1	0.1

Summary

- Ground truth is never known $\Rightarrow$ replace by report of a random peer.
- Constant reward: report most common answer.
- Peer prediction: use proper scoring rule to scale rewards: report truthfully even uncommon answers.
- However, *posterior* distributions needs to be common and known!
- Peer Truth Serum: *prior* distribution needs to be common and known, updates need to satisfy *self-predicting* condition.
- Peer Truth Serum for Crowdsourcing: no need for common prior, but requires set of similar tasks.

References

R. Jurca and B. Faltings: *An incentive compatible reputation mechanism*, IEEE International Conference on e-commerce. IEEE, 2003. pp. 285-292, 2003.

N. Miller, P. Resnick, R. Zeckhauser: *Eliciting informative feedback: The peer-prediction method*, Management Science **51**(9), pp. 1359-1373, 2005.

L. Von Ahn, L. Dabbish: *Designing games with a purpose*, Communications of the ACM **51**(8) pp. 58-67, 2008.

R. Jurca and B. Faltings: *Mechanisms for making crowds truthful*, Journal of Artificial Intelligence Research **34**(1) pp. 209-253, 2009

R. Jurca and B. Faltings: *Incentives for answering hypothetical questions*, Workshop on Social Computing and User Generated Content, EC-11, 2011.

J. Witkowski and D. Parkes: *Peer prediction without a common prior*, Proceedings of the 13th ACM Conference on Electronic Commerce, pp. 964-981, 2012.

S.-W. Huang and W.-T. Fu: *Enhancing reliability using peer consistency evaluation in human computation*, Proceedings of the 2013 conference on Computer supported cooperative work, pp. 639-648, 2013

G.Radanovic et al.: *Incentives for Effort in Crowdsourcing using the Peer Truth Serum*, ACM TIST, 2016