Algorithmic Bayesian persuasion with combinatorial actions

Kaito Fujii (National Institute of Informatics)

(joint work with **Shinsaku Sakaue** (UTokyo))

AAAI 2022

A game-theoretic model of strategic information revelation for how to lead an agent to a preferred action

Bayesian persuasion [Kamenica–Gentzkow'11]

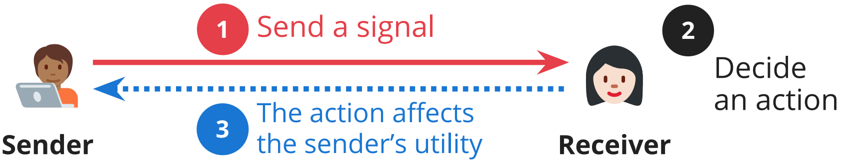

Q **How to compute an optimal signaling strategy?**

Applications of Bayesian persuasion

(1) Ads

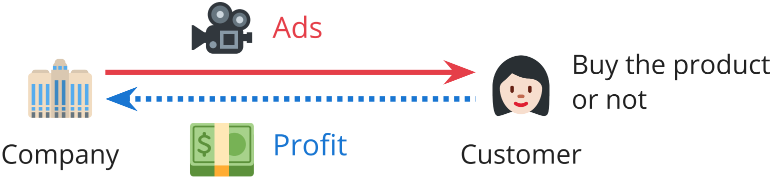

(2) Election

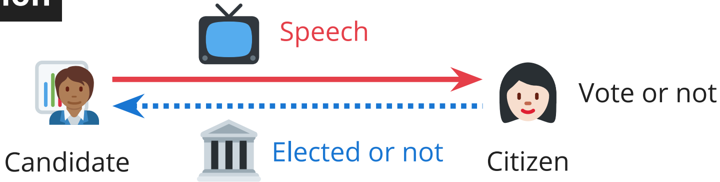

Example of Bayesian persuasion [Kamenica–Gentzkow'11]

Q **What is an optimal signaling strategy for the company?**

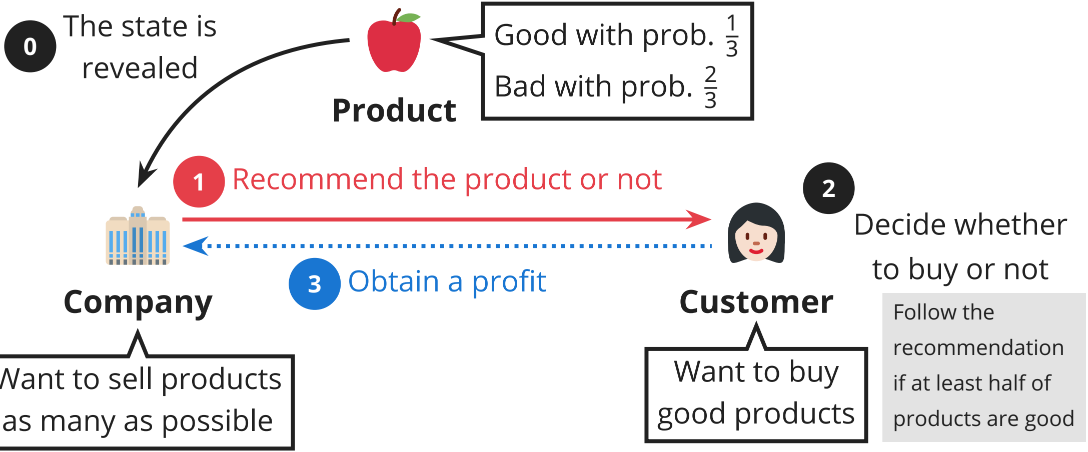

Example of Bayesian persuasion [Kamenica–Gentzkow'11]

Optimal Recommend all good products and half of bad products

→ Customer buys $\frac{2}{3}$ of all products

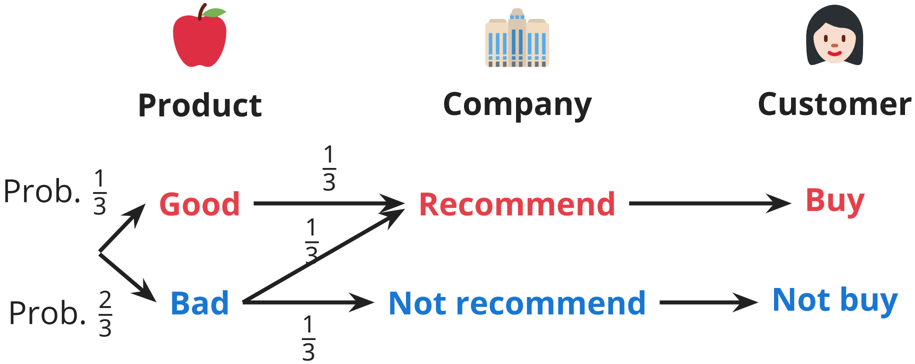

Example of Bayesian persuasion [Kamenica–Gentzkow'11]

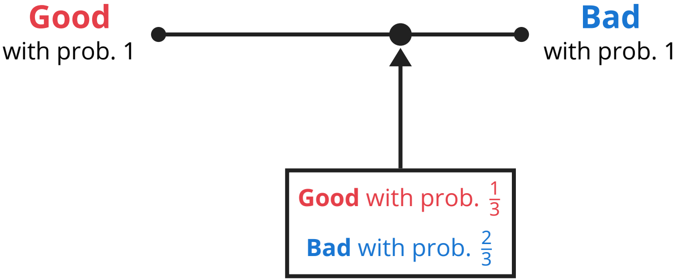

Example of Bayesian persuasion [Kamenica–Gentzkow'11]

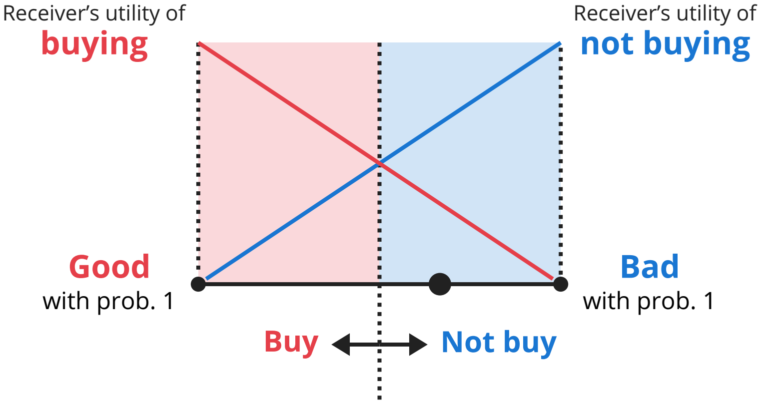

Notation for Bayesian persuasion

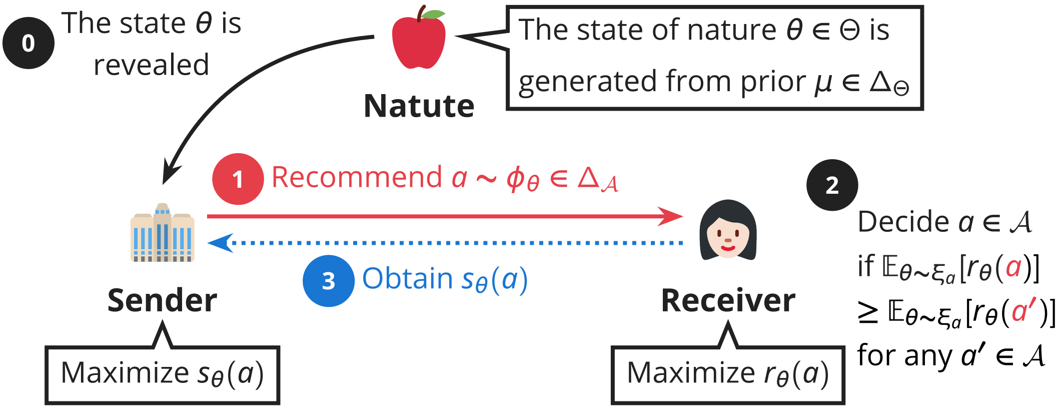

※ signaling scheme $(\phi_\theta)_{\theta\in\Theta}$ is declared in advance

(commitment assumption)

Linear programming formulation [Kamenica–Gentzkow'11]

Computation of an optimal strategy can be formulated as LP

$$\underset{(\phi_\theta)_{\theta\in\Theta}}{\text{maximize}} \quad \sum_{\theta\in\Theta}\sum_{a\in\mathcal{A}} \mu(\theta)\phi_\theta(a)s_\theta(a)$$

$$= \mathbb{E}_{\theta\sim\mu, a\sim\phi_\theta}[s_\theta(a)]$$ **Sender's expected utility**

$$\text{subject to} \quad \sum_{\theta\in\Theta} \mu(\theta)\phi_\theta(a)\big(r_\theta(a) - r_\theta(a')\big) \geq 0 \quad (a, a' \in \mathcal{A})$$

$$\iff \mathbb{E}_{\theta\sim\xi_a}[r_\theta(a)] \geq \mathbb{E}_{\theta\sim\xi_a}[r_\theta(a')]$$ **persuasiveness constraints**

$$\phi_\theta \in \Delta_{\mathcal{A}} \quad (\theta \in \Theta)$$

Bayesian persuasion with combinatorial actions

(1) Buy multiple products

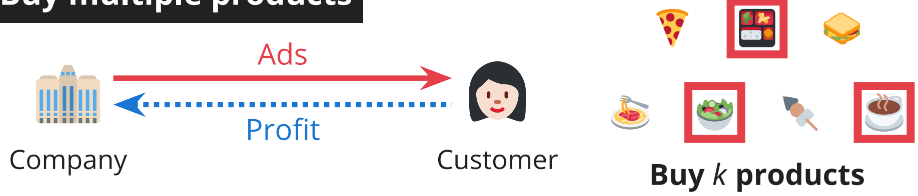

(2) Choose a path

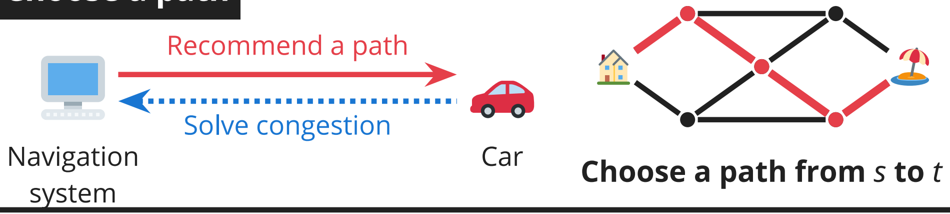

Bayesian persuasion with combinatorial actions

Receiver's action is a combination of elements in a finite set E

i.e., $\mathcal{A} = \mathcal{I}$, where $\mathcal{I} \subseteq 2^E$ is a set family

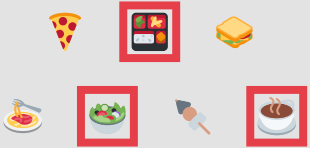

Buy k products

E is the set of products

$\mathcal{I} = \{S \subseteq E \mid |S| \leq k\}$

uniform matroid constraints

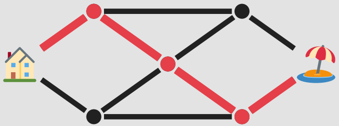

Choose a path from s to t

E is the set of edges

$\mathcal{I} = \{S \subseteq E \mid s\text{–}t \text{ path}\}$

path constraints

Notation for BP with combinatorial actions

Consider a recommendation of combinatorial action $S \in \mathcal{I} \subseteq 2^E$

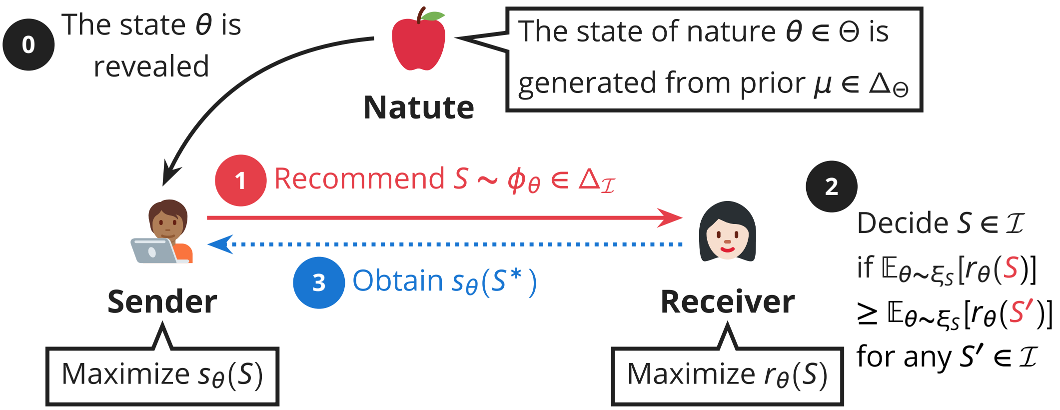

LP formulation of BP with combinatorial actions

LP with exponentially many variables and constraints

$$\underset{(\phi_\theta)_{\theta\in\Theta}}{\text{maximize}} \quad \sum_{\theta\in\Theta}\sum_{S\in\mathcal{I}} \mu(\theta)\phi_\theta(S)s_\theta(S) \qquad \textbf{Sender's expected utility}$$

$$\text{subject to} \quad \sum_{\theta\in\Theta} \mu(\theta)\phi_\theta(S)\big(r_\theta(S) - r_\theta(S')\big) \geq 0 \quad (S, S' \in \mathcal{I})$$

persuasiveness constraints

$$\phi_\theta \in \Delta_{\mathcal{I}} \quad (\theta \in \Theta)$$

Q Is it possible to solve this LP in time polynomial in $|E|$?

Summary of our contributions

Q Is it possible to efficiently compute an optimal signaling strategy for Bayesian persuasion with combinatorial actions?

Our results

1. **NP-hardness** of constant-factor approximations
2. Poly-time algorithms when **the number of states is a constant**
3. Poly-time algorithms for **CCE-persuasiveness**

Result (1): NP-hardness

Constant-factor approximation is NP-hard for simple constraints

Theorem

For any $\alpha \in (0, 1]$, it is NP-hard to compute an α-approximate signaling scheme for Bayesian persuasion with any of

- uniform matroid constraints
- partition matroid constraints
- graphic matroids constraints
- path constraints

if the utility functions are linear, i.e., $s_\theta(S) = \sum_{i \in S} s_\theta(\{i\})$ and $r_\theta(S) = \sum_{i \in S} r_\theta(\{i\})$

Result (1): NP-hardness

- **Partition matroids**

  Reduction from **public Bayesian persuasion with no externalities**
  [Dughmi–Xu'17]

- **Uniform matroids, Graphic matroids, Paths**

  Reduction from **LINEQ-MA**$(1-\zeta, \delta)$ [Guruswami–Raghavendra'09]

  Given a linear system $Ax = c$, the promise problem of distinguishing

  – there exists $x \in \{0, 1\}^n$ that satisfies at least a $1-\zeta$ fraction of the equations

  – every $x \in \mathbb{Q}^n$ satisfies less than a δ fraction of the equations

  based on the reduction for OptSignal [Castiglioni–Celli–Marchesi–Gatti'20]

Result (2): When $|\Theta|$ is a constant

We need to consider combinations that can be a best response

$$\sum_{\theta\in\Theta} \mu(\theta)\phi_\theta(S)\,(r_\theta(S) - r_\theta(S')) \geq 0 \quad (S, S' \in \mathcal{I})$$

persuasiveness constraints

$$\Leftrightarrow S \in \operatorname*{argmax}_{S\in\mathcal{I}} \sum_{\theta\in\Theta} \mu(\theta)\phi_\theta(S) r_\theta(S) \text{ for each } S \in \mathcal{I}$$

S is a best response for posterior $\xi_S(\theta) \propto \mu(\theta)\phi_\theta(S)$

Observation

In the LP formulation, instead of $\mathcal{I}$, it is sufficient to consider

$$\mathcal{I}^* = \left\{ S \in \mathcal{I} \;\middle|\; \exists \xi \in \Delta_\Theta : S \in \operatorname*{argmax}_{S\in\mathcal{I}} \sum_{\theta\in\Theta} \xi(\theta) r_\theta(S) \right\}$$

Result (2): When $|\Theta|$ is a constant

Assume $r_\theta(\cdot)$ is linear, i.e., $r_\theta(S) = \sum_{i \in S} r_\theta(\{i\})$ and $s_\theta(\cdot)$ is given by a value oracle

Uniform matroids

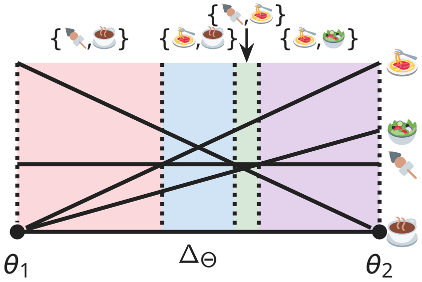

For each posterior $\xi \in \Delta_\Theta$,
selecting top-k elements is best
↓
k-level faces enumeration [Mulmuley'91]
studied in computational geometry

Lemma

Under a certain degeneracy assumption, $|\mathcal{I}^*| = O(|E|^{|\Theta|-1})$

General matroids

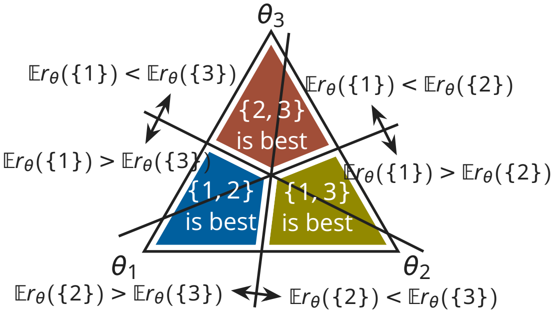

For each posterior $\xi \in \Delta_\Theta$,

the ranking of $\mathbb{E}r_\theta(\{\cdot\})$ determines the best response

↓

cell enumeration in an arrangement of hyperplanes [Edelsbrunner'87]

Lemma

Under a certain degeneracy assumption, $|\mathcal{I}^*| = O(|E|^{2(|\Theta|-1)})$

Result (2): When $|\Theta|$ is a constant

- **Partition matroids**

  It is sufficient to consider the ranking in each partition

  ➡ $|\mathcal{I}^*| = O(|E|^{2(|\Theta|-1)}/P^{(|\Theta|-1)})$, where P is the number of partitions

- **Graphic matroids**

  The enumeration is reduced to the parametric spanning tree problem

  $|\mathcal{I}^*| = O(|E||V|^{1/3})$ when $|\Theta| = 2$, where V is the set of vertices [Dey'98]

- **Paths**

  Even if $|\Theta| = 2$, there exists an instance such that $|\mathcal{I}^*| = |E|^{\Omega(\log |E|)}$

  [Carstensen'83]

Result (3): CCE-persuasiveness

In the CCE-persuasiveness setting, the receiver selects either of:

- following the signal $\left(\text{the expected utility is } \sum_{\theta \in \Theta} \sum_{S \in \mathcal{I}} \mu(\theta)\phi_\theta(S) r_\theta(S)\right)$
- not receiving the signal $\left(\text{the expected utility is } \max_{S \in \mathcal{I}} \sum_{\theta \in \Theta} \mu(\theta) r_\theta(S)\right)$

$$
\begin{aligned}
\text{maximize} \quad & \sum_{\theta \in \Theta} \sum_{S \in \mathcal{I}} \mu(\theta)\phi_\theta(S) s_\theta(S) \\
\text{subject to} \quad & \sum_{\theta \in \Theta} \sum_{S \in \mathcal{I}} \mu(\theta)\phi_\theta(S) r_\theta(S) \geq \max_{S \in \mathcal{I}} \sum_{\theta \in \Theta} \mu(\theta) r_\theta(S) \\
& \phi_\theta \in \Delta_{\mathcal{I}} \quad (\theta \in \Theta)
\end{aligned}
$$

CCE-persuasiveness constraints

Result (3): CCE-persuasiveness

Theorem (informal)

If we have an oracle that, for any $y \geq 0$ and $\theta \in \Theta$, returns $S \in \mathcal{I}$ s.t.

$$s_\theta(S) + y \cdot r_\theta(S) \geq \max_{S' \in \mathcal{I}} \alpha \cdot s_\theta(S') + y \cdot r_\theta(S'),$$

then we can compute an $(\alpha - \epsilon)$-approximation for any $\epsilon \in (0, \alpha)$

Proof Consider a separation oracle for the dual LP

Applications

- s_θ, r_θ: linear, $\mathcal{I}$: matroid, $\alpha = 1$
- s_θ: monotone submodular, r_θ: linear, $\mathcal{I}$: matroid, $\alpha = 1 - 1/e$

Summary of our contributions

Q Is it possible to efficiently compute an optimal signaling strategy for Bayesian persuasion with combinatorial actions?

Our results

1. **NP-hardness** of constant-factor approximations
2. Poly-time algorithms when **the number of states is a constant**
3. Poly-time algorithms for **CCE-persuasiveness**

Reference

- Matteo Castiglioni, Andrea Celli, Alberto Marchesi, and Nicola Gatti: Online Bayesian Persuasion. In *NeurIPS 2020*.
- P. J. Carstensen: Complexity of some parametric integer and network programming problems. *Mathematical Programming* 26(1): 64–75 (1983).
- T. K. Dey: Improved bounds for planar k-sets and related problems. *Discrete & Computational Geometry* 19(3): 373–382 (1998).
- Shaddin Dughmi and Haifeng Xu: Algorithmic persuasion with no externalities. In *EC 2017*.
- H. Edelsbrunner: *Algorithms in Combinatorial Geometry* (1987).
- Emir Kamenica and Matthew Gentzkow: Bayesian persuasion. *American Economic Review* 101(6): 2590–2615 (2011).
- Ketan Mulmuley: On levels in arrangements and Voronoi diagrams. *Discrete & Computational Geometry* 6(3): 307–338 (1991).
- Venkatesan Guruswami and Prasad Raghavendra: Hardness of Learning Halfspaces with Noise. *SIAM Journal on Computing* 39(2): 742–765 (2009).