Optimized Distortion In Linear Social Choice

Luise Ge, Gregory Kehne, Yevgeniy Vorobeychik
Washington University in St. Louis
@ AAAI 2026

Single-Winner Selection

We have a set of n voters: V

and a set of m candidates: C

Given a ranking $\vec{\sigma}$:

$$v_1 : c_1 \succ c_2 \succ c_3 \ldots$$
$$v_2 : c_2 \succ c_3 \succ c_1 \ldots$$

How to compute $f(\vec{\sigma}) \in \Delta(C)$?

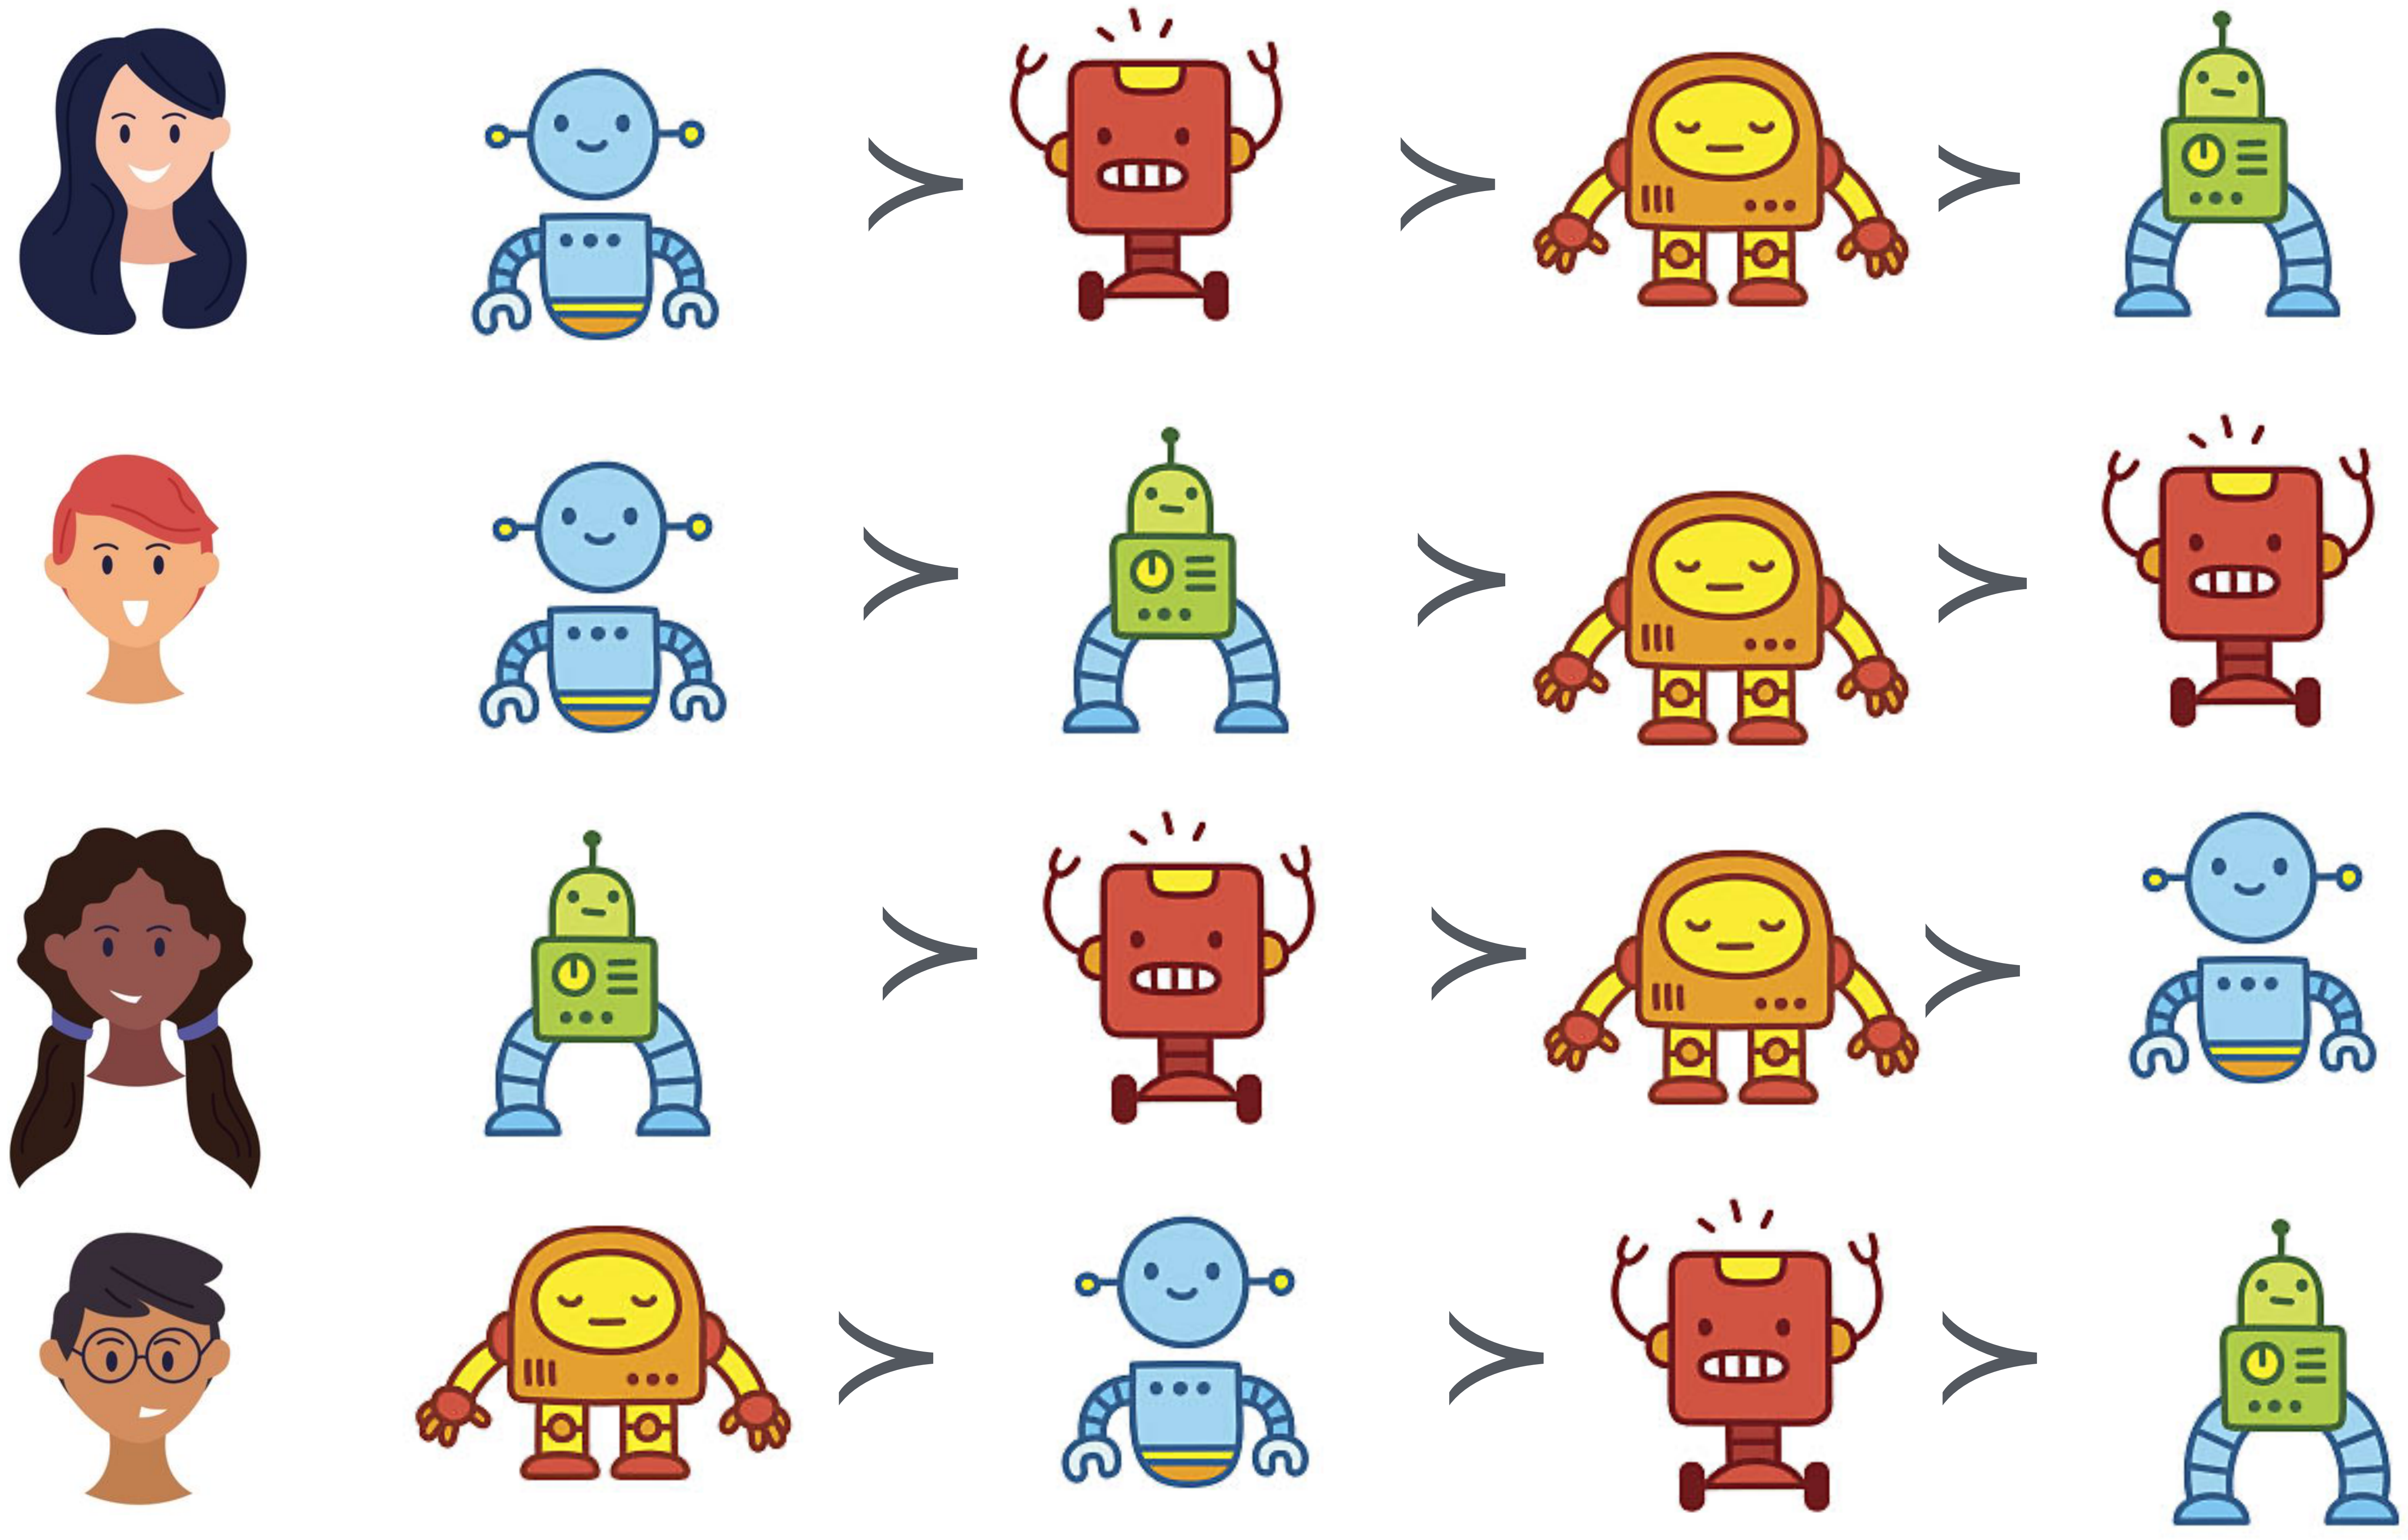

Distortion: ordinal v.s. cardinal

Suppose utilites exist and $c \succ_v c'$ ***iff*** $u_v(c) \geq u_v(c')$. *How to maximize*

$$\mathrm{UW}(\vec{\mathrm{u}}, \mathrm{c}) = \sum_{v \in V} u_v(c)?$$

	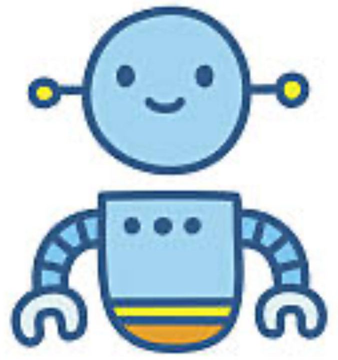	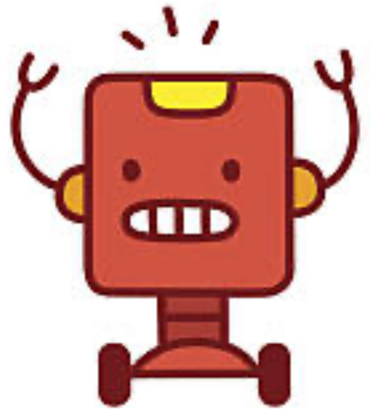	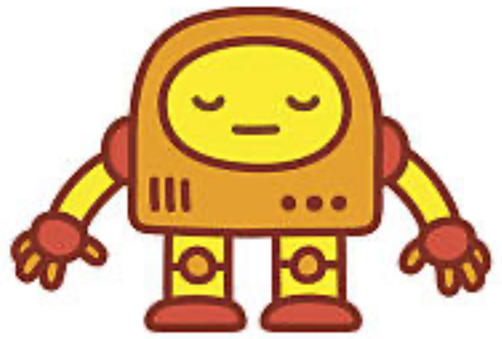	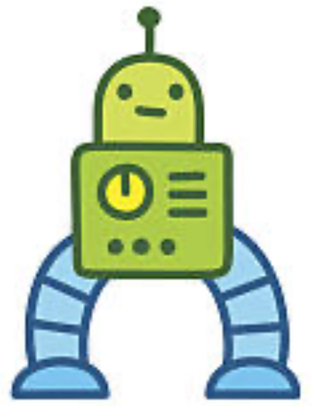
	0.4	0.3	0.3	0
	0.4	0.3	0	0.3
	0.25	0.25	0.25	0.25
	0.1	0.8	0.05	0.05
UW	1.15	1.65	0.6	0.6

As there are many utility profiles consistent with the rankings, we want a worst-case guarantee:

$$D(f, \vec{\sigma}) := \max_{\vec{u} \triangleright \vec{\sigma}} \frac{\mathrm{UW}(\vec{u}, c^*(\vec{u}))}{\mathbb{E}_{c \sim f(\vec{\sigma})}[\mathrm{UW}(\vec{u}, c)]}$$

$$D(f) := max_{\vec{u} \triangleright \vec{\sigma}, \vec{u} \in \mathcal{U}^n} D(f, \vec{\sigma})$$

[Procaccia and Rosenschein 2006]

Classical model(s)

- The denominator should be nonzero (No-mixed signs for utilities)
- No one should dominate (Normalization)

Most common one (Unit-Sum):

Each voter has her utilities for candidates summing to one.

Best deterministic rule has distortion $\Theta(m^2)$ [Caragiannis et al. 2015]

Best randomized distortion is $\Theta(\sqrt{m})$ [Ebadian et al. 2024]

What if m is *unbounded* ??

Linear Social Choice

Axioms for AI Alignment From Human Feedback [Ge et. al. 2024]

Each voter v and each candidate c has an embedding in $\mathbb{R}^d$ and the utility is linear: $u_v(c) = v^T c$

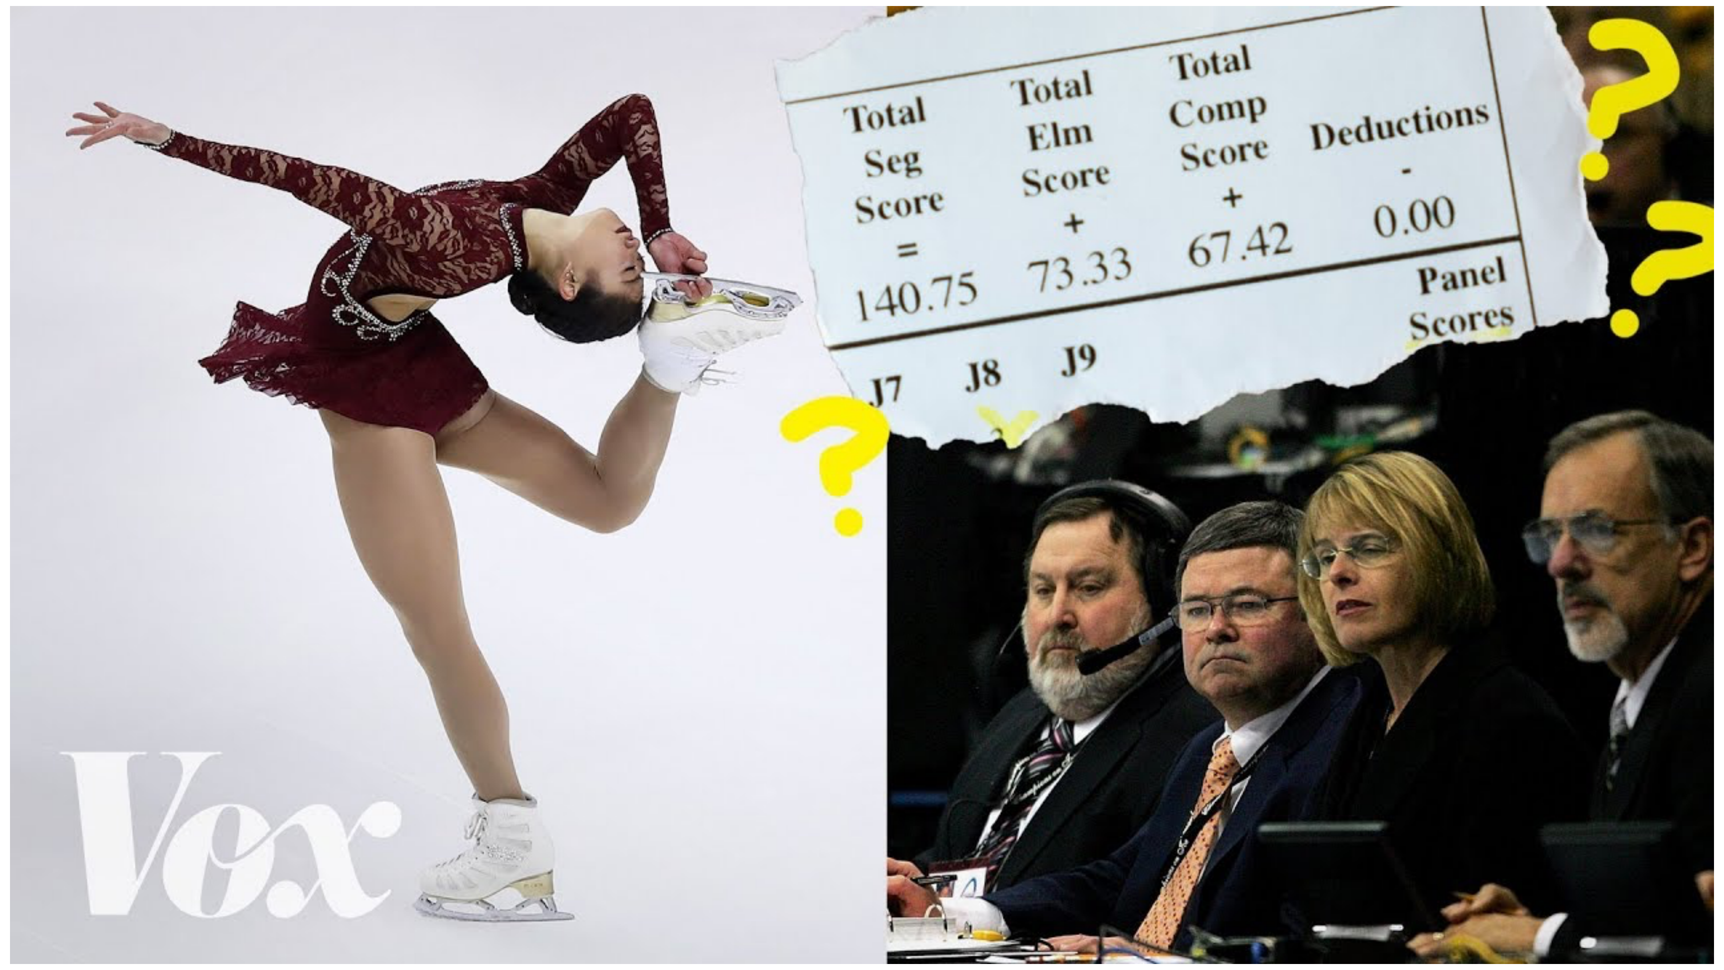

Multi-objective Evaluation

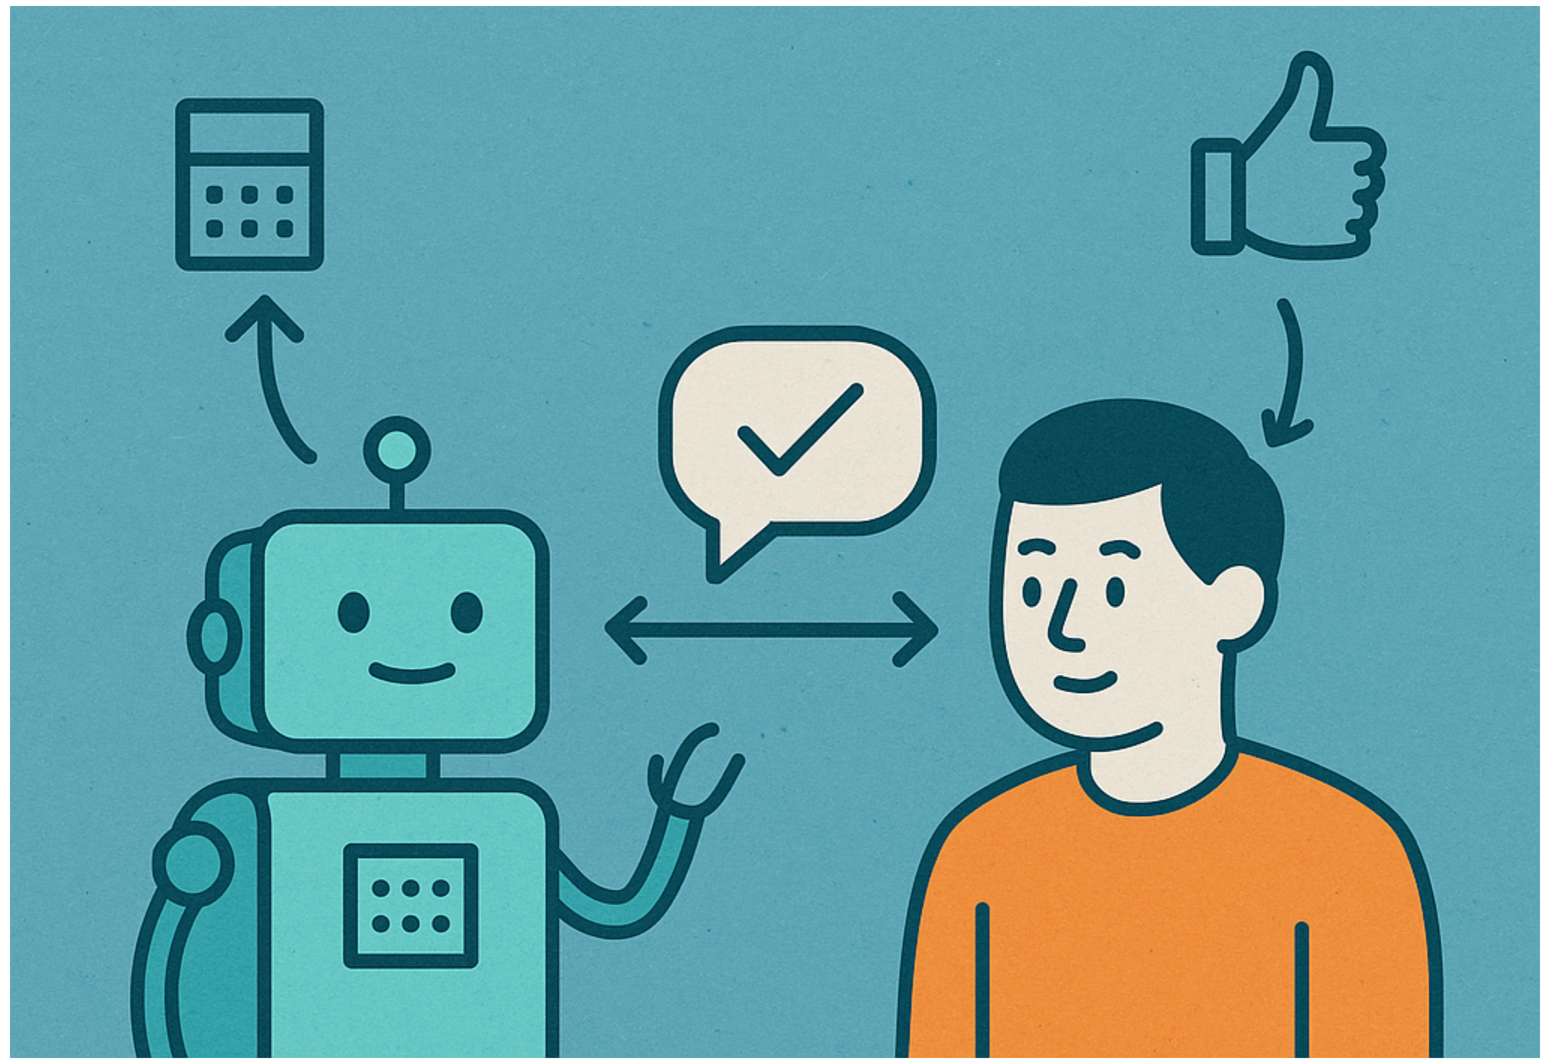

Reinforcement Learning from Human Feedback (RLHF):

$r(\text{prompt}, \text{answer}) = \theta^T \phi(\text{prompt}, \text{answer})$

Collaborative filtering

$U = \tilde{V}\tilde{C}^T$

Model 1: ℓ_1 normalized

- The denominator should be nonzero (No-mixed signs for utilities)
- No one should dominate (Normalization)

$$v \in \Delta_d, \quad c \in \Delta_d$$

$$V \in Cone(C)$$

$$u_v(c) = v^T c$$

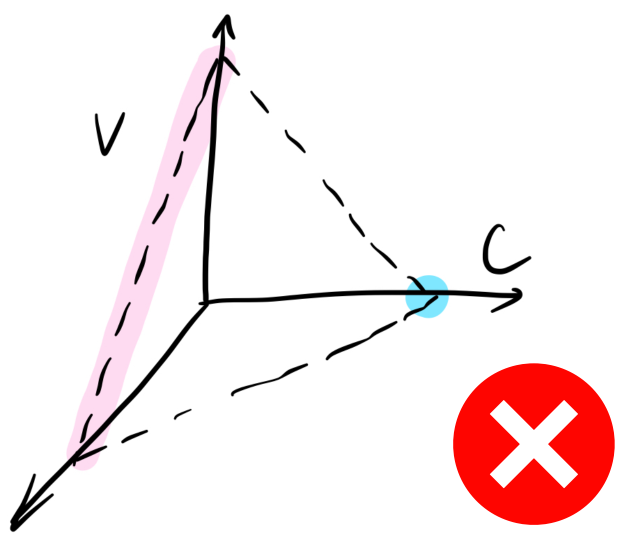

Observe that the unit-sum assumption becomes a special case! Consider $d = m$*, and* $c_i = e_i$ *for* $i \in [m]$.

Model 2: ℓ_2 normalized

$$v \in \mathbb{S}^d_+, \quad c \in \mathbb{S}^d_+$$

$$u_v(c) = v^T c = cos(\angle(v, c))$$

$$V \in Cone(C)$$

Then this is related to the metric distortion if we consider the distance metric $d(v, c)$ *to be:* $1 - u_v(c)$*, the cosine distance.*

Results with L1-normalization

Unit-Sum

Plurality: $O(m^2)$

Deterministic rules: $\Omega(m^2)$

Best randomized rule: $\Theta(\sqrt{m})$

Linear Social Choice

selects the plurality winner among the set: $\hat{C} := \left\{ c_i \in C \mid c_i = \arg\max_{c \in C} c^i \text{ for each } i \in [d] \right\}$

Plurality: $O(\min(m, n)d)$

Maximum Coordinate Plurality: $O(d^3)$

Deterministic rules: $\Omega(d^2)$

Best randomized distortion is $\Theta(\sqrt{d})$

Linear Stable Lottery Rule: half-chance the reverse-KL projection to the center vector 1/d; half-chance stable lottery (*Optimized Distortion and Proportional Fairness in Voting* [Ebadian et al. 2024])

Takeaway: if further utiliy structure applies, the distortion bounds depend more on the embeddings' complexities than the number of candidates.

Instance-Optimal Rule

For any fixed ranking, we can efficiently compute a determinisitc/ randomized voting rule in **poly(m,n,d)** time, whose distortion is near-optimal for that instance.

Future Work

- **In general, linear social choice is an exciting area with lots to be done!!**
  - Rules with guarantees on other objectives (Nash Product, Regret...)
  - Multi-Winner Elections
  - More empirical studies, even extend it to nonlinear utility functions.

- **Specific to our paper**
  - Close the gaps between upper and lower bounds, for instance, deterministic rules has bounds $\Omega(d^2), O(d^3)$
  - What bound does the instance-optimal rule have?

Thank you!

All these bounds depend only on d, great when $d \ll m$!

Maximum Coordinate Plurality (MCP): $O(d^3)$

It selects the plurality winner among the set $\hat{C} := \left\{ c_i \in C \mid c_i = \arg\max_{c \in C} c^i \text{ for each } i \in [d] \right\}$

Uniform Projection (UProj): $O(d)$

It defines a distribution such that the expected candidate vector corresponds to

$$\hat{c} := \arg \min_{x \in Conv(C)} \mathrm{KL}(\mu \| x), \quad \mu = \overrightarrow{\frac{1}{d}}.$$

Linear Stable Lottery (LSLR): $O(\sqrt{d})$

Given a stable lottery $\mathscr{W}$ over committees of size $\sqrt{d}$, the it chooses each c with probability

$$\frac{1}{2\sqrt{d}} \Pr_{W \sim \mathscr{W}(\vec{\sigma})}[c \in W] + \frac{1}{2} \Pr_{c' \sim f_{\mathrm{Uproj}}(\vec{\sigma})}[c' = c].$$

Definition 3.1. A distribution $\mathbf{X} \in \Delta(\mathcal{P}_k(A))$ over committees X of size k is a *stable lottery* if for all alternatives $a^* \in A$, we have

$$\mathbb{E}_{X \sim \mathbf{X}} [V(a^*, X)] < \frac{n}{k}.$$