TO LEAD OR TO BE LED:

A GENERALIZED CONDORCET JURY THEOREM UNDER DEPENDENCE

Jonas Karge, Juliette-Michelle Burkhardt, Sebastian Rudolph, Dominik Rusovac

Computational Logic Group, TU Dresden

AAMAS 2024, Auckland, New Zealand, May 08, 2024

The Condorcet Jury Theorem (CJT)

Marie Jean Antoine Nicolas Caritat Marquis de Condorcet

Theorem: For odd-numbered **homogenous** groups of **independent** and **reliable** agents in a **dichotomic** voting setting, the probability that majority voting identifies the correct alternative

- increases monotonically with the number of agents and (non-asymptotic part)
- converges to 1 as the number of agents goes to infinity. (asymptotic part)

Past Generalizations

Weakening Voting Constraints.

- Weakened dichotomy: Any finite number of alternatives (List and Goodin 2001).
- Weakened completeness: Vote for any number of alternatives (Everaere, Konieczny, and Marquis 2010).

Weakening Homogeneity.

- Independence, completeness, dichotomy + heterogeneous competence levels (Owen, Grofman, and Feld 1989).
- Indepedence + heterogenous competence levels (Karge and Rudolph 2022).

Weakening Independence with Opinion Leader.

- Dichotomy, completeness, homogeneity + equally competent OL (Boland, Proschan, and Tong 1989).
- Dichotomy, completeness, homogeneity + OL with deviating competency (Goodin and Spiekermann 2018).

Goal: Weaken all assumptions simultaneously.

Voting

Define **approval voting** and obtain simpler voting mechanisms as special cases.

Given: finite set of n agents $\mathcal{A} = \{a_1, \ldots, a_n\}$
finite set of m choices $\mathcal{W} = \{\omega_1, \ldots, \omega_m\}$

- **approval voting (instance)**: relation $V \subseteq \mathcal{A} \times \mathcal{W}$
  $(a_i, \omega_j) \in V$ means agent a_i approves choice ω_j
- given $\omega \in \mathcal{W}$, obtain **score** $\#_V\, \omega$ as overall number of votes that ω receives, i.e.,

$$\#_V\, \omega = |\{a_i \in \mathcal{A}_n \mid (a_i, \omega) \in V\}|$$

- ω **wins approval vote** V if it receives strictly more votes than any other choice:

$$\#_V\, \omega > \max_{\omega' \in \mathcal{W} \setminus \{\omega\}} \#_V\, \omega'$$

The Probabilistic Framework

Random process chooses ω_* (the actual world state) and generates V, governed by **joint probability distribution** $\mathbb{P}$ over Bernoulli (i.e., $\{0, 1\}$-valued) random variables

$$
\begin{array}{ccc}
X_*^{\omega_1}, & \dots, & X_*^{\omega_m}, \\
X_o^{\omega_1}, & \dots, & X_o^{\omega_m}, \\
X_1^{\omega_1}, & \dots, & X_1^{\omega_m}, \\
\vdots & \vdots\vdots\vdots & \vdots \\
X_n^{\omega_1}, & \dots, & X_n^{\omega_m}.
\end{array}
$$

- $X_*^{\omega_j}$ is 1 if ω_j is the true world state (i.e., $\omega_j = \omega_*$), else 0,
- $X_o^{\omega_j}$ is 1 if the OL approves ω_j, and 0 otherwise,
- $X_i^{\omega_j}$ represents the private signal of the ith agent regarding his approval of the jth world state: it is 1 if a_i privately approves ω_j and otherwise 0.

The Joint Probability Distribution – Assumptions

Definition: A joint distribution satisfies **private agent approval independence** if, conditioned on the actual world state, the private decision to approve any given ω_j is made independently across all agents, i.e., for any $\omega, \omega_j \in \mathcal{W}$ and any sequence $v_1, \ldots, v_n$ of values from $\{0,1\}$ the following holds:

$$\mathbb{P}\Big(\bigwedge_{i=1}^{n} X_i^{\omega_j} = v_i \mid [\omega_* {=} \omega]\Big) = \prod_{i=1}^{n} \mathbb{P}\Big(X_i^{\omega_j} = v_i \mid [\omega_* {=} \omega]\Big).$$

Let p_k^{ω} denote an agents internal competency and $\bar{p}^{\omega} = \frac{1}{n} \sum_{k=1}^{n} p_k^{\omega}$.

Definition: A joint probability distribution satisfies **Δp-group reliability** for some $\Delta p > 0$, if for every n and $\omega_\dagger \in \mathcal{W} \setminus \{\omega_*\}$ holds

$$\bar{p}^{\omega_*} \geq \Delta p + \bar{p}^{\omega_\dagger}.$$

The Voting Scenario

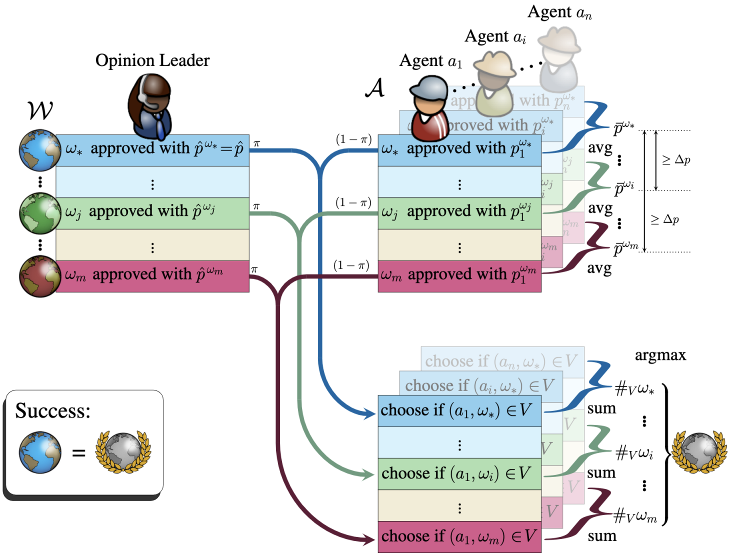

Asymptotic Result

Theorem: Taking into account the OL's competency $\hat{p}$, we obtain the following probability that converges to 1 with growing n:

$$\mathbb{P}\Big(\bigwedge_{\omega_\dagger \in \mathcal{W}\setminus\{\omega_*\}} V^{\omega_*} > V^{\omega_\dagger}\Big) \geq 1-(m-1)\Big(\hat{p}e^{-\frac{1}{2}n\Delta p^2(1-\pi)^2} + (1-\hat{p})e^{-\frac{1}{2}n(\Delta p(1-\pi)-\pi)^2}\Big).$$

Proof Idea.

- Distinguish case where OL is right/wrong.
- Let $\omega_\dagger \in \mathcal{W} \setminus \{\omega_*\}$ denote an arbitrary but fixed "competitor" of ω_* in the approval vote.
- For each case, apply **Hoeffdings's inequality** to obtain lower bound for the probability of ω_* winning against $\omega_\dagger$.
- Obtain the probability for ω_* winning the approval vote against **all** competing $\omega_\dagger \in \mathcal{W} \setminus \{\omega_*\}$ simultaneously.
- Combine both bounds through $\hat{p}$.

Important Caveat

When the OL is wrong, π may not be too large.

Threshold: We require that

$$\pi < \frac{\Delta p}{\Delta p + 1}.$$

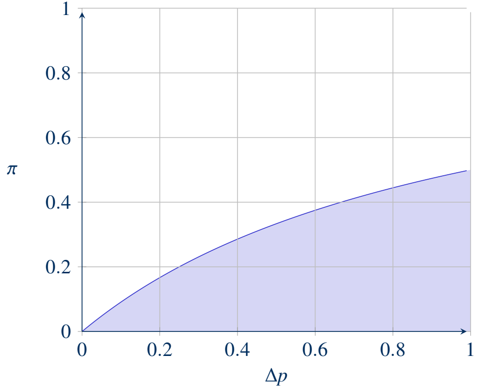

Estimates for Required Number of Agents

Theorem allows to derive bound on number n of agents required for success with probability of at least $P_{\min}$, given

- $\Delta p \in (0, 1]$
- $m > 1$
- $P_{min} < 1$
- $\pi \in [0, \frac{\Delta p}{\Delta p+1})$
- $\hat{p} \in [0, 1]$

Theorem:

$$\hat{p}e^{-\frac{1}{2}n\Delta p^2(1-\pi)^2} + (1-\hat{p})e^{-\frac{1}{2}n(\Delta p(1-\pi)-\pi)^2} \leq \frac{1 - P_{\min}}{m-1} \tag{1}$$

$$\frac{2}{(\Delta p(1-\pi) - \pi)^2} \ln \frac{m-1}{1 - P_{\min}} \leq n \tag{2}$$

Experimental Design

Goal: Test tightness of our bounds.

- Reproduce worst-case circumstances (homogeneous competencies, OL approves all incorrect choices).
- High $P_{min} = 0.9$, simple setting with $m = 2$, varying $\hat{p}$.
- Iterated through different Δp-values.
- Tested values for n by simulating high number of voting rounds to find number needed to surpass $P_{min} = 0.9$.
- Identification of n sped up with binary stochastic search.

Fixing Δp, varying $\hat{p}$

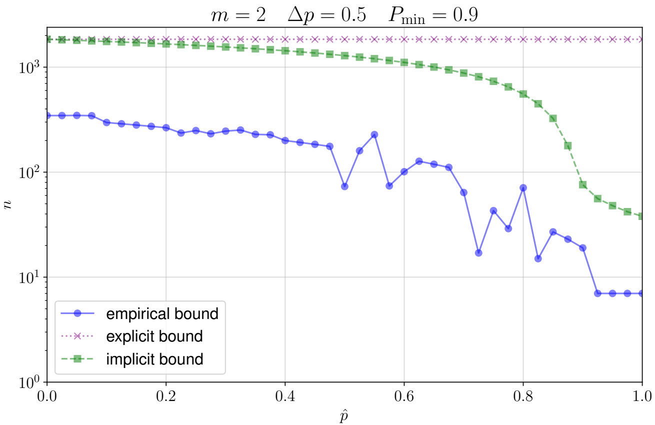

Figure: Influence of $\hat{p}$ on the empirical and bound values.

Summary

Generalized the CJT to a setting that allows

- heterogeneous competence levels among agents;
- approval voting for any (finite) number of alternatives.
- dependence among voters through an OL.

For this setting, we

- determined the amount of influence the OL may exert;
- derived practical estimates for the number of independent agents necessary to guarantee a prescribed minimal probability of success;
- examined the bound's tightness by means of statistical simulations.