TECHNISCHE UNIVERSITÄT DRESDEN

SECAI

ELICIT AND WEIGH: A VOTING-BASED APPROACH TO OPTIMAL WEIGHTS IN IMPRECISE LINEAR POOLING

Lea Bauer and Jonas Karge

Computational Logic Group

School of Embedded Composite Artificial Intelligence (SECAI)

ECSQARU, Hagen, September 25th, 2025

High Level View

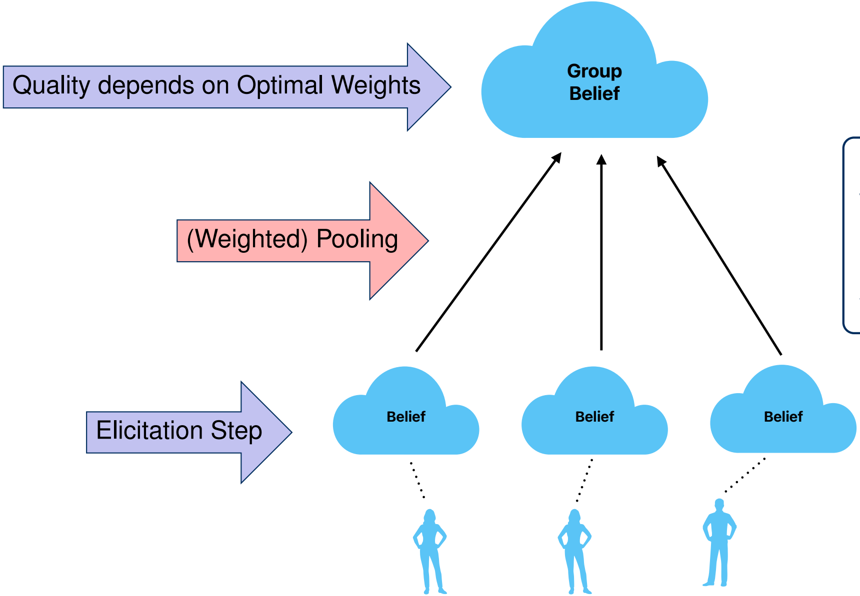

Idea: Derive optimal weights by modeling the elicitation step as a voting problem.

Imprecise Opinion Pooling

Imprecise Opinion Pooling

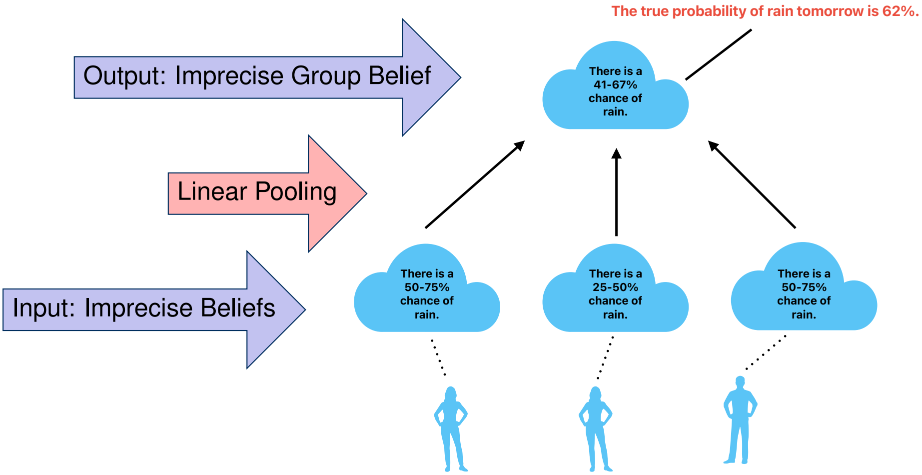

Imprecise Pooling

Scenario: Multiple experts assess the likelihood of an event such as:

> **Example:** It will rain in Hagen on Monday of next week.

$\Rightarrow$ We model the probabilistic beliefs $\mathcal{P}_i(A)$ that agent i holds about a proposition A as intervals of probability values of the form $\mathcal{P}_i(A) = [a, b]$.

An imprecise pooling function takes as **input** n imprecise beliefs, one for each agent, for an event and yields as **output** a single collective imprecise belief.

> **Definition: Linear Pooling.** $\mathcal{F}([a_1, b_1], \ldots, [a_n, b_n])(A) = [\sum_i \lambda_i a_i, \sum_i \lambda_i b_i]$.

The input profile is defined in terms of the lower and upper probabilities where λ_i denote the weight assigned to agent i's belief.

Imprecise Linear Pooling - Weights

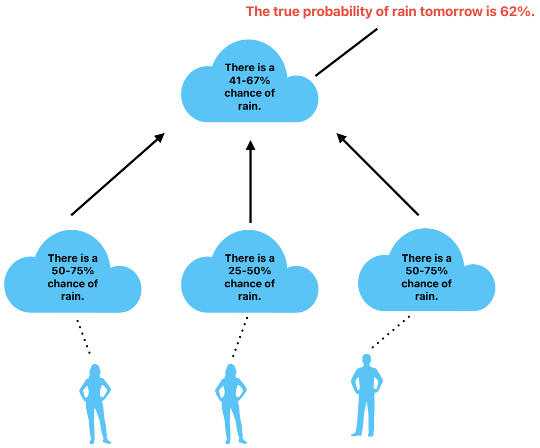

Imprecise Linear Pooling - Weights

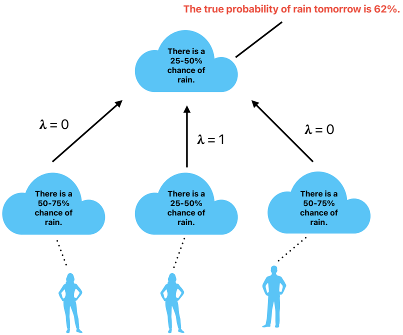

Imprecise Linear Pooling - Weights

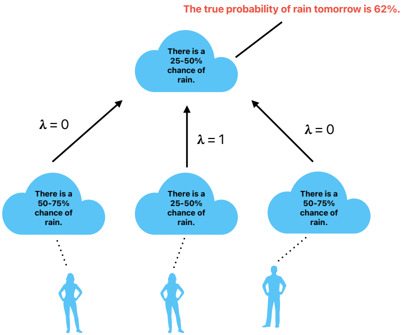

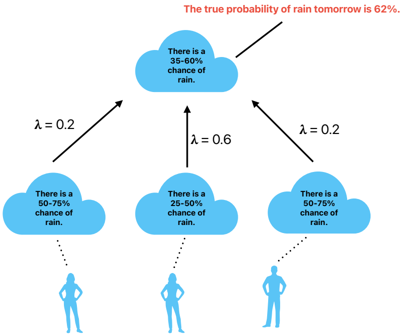

Belief Elicitation through Voting

Epistemic Voting

Suppose, we are dealing with

- a set of **agents** (people, sensors, drones, ...)
- that **vote** (via some voting rule)
- for **alternatives** (policies, interpretations of sensor data, courses of action, ...).

Two distinct goals for voting procedures:

(1) Ensure a fair voting procedure;

(2) identify the correct alternative.

We assume: There is exactly one correct alternative, the **ground truth**.

Belief Elicitation

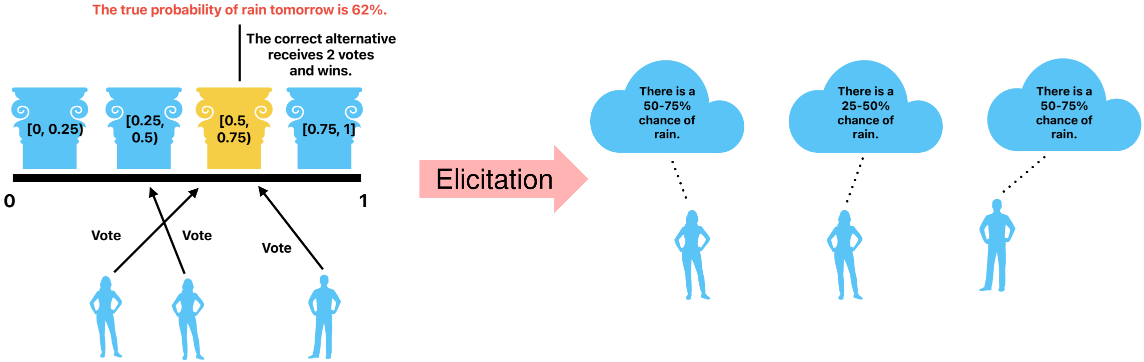

Elicitation through Plurality Voting

We associate each bin with an alternative ω_j in the voting process:

Definition: Bin. Each alternative $\omega_k \in \mathcal{W} = \{\omega_1, \dots, \omega_m\}$ represents a subinterval (bin) of the form $[a, b)$, obtained by partitioning the unit interval such that each ω_k is of equal size $l := (b - a)$. The final subinterval is of the form $[a_{final}, 1]$.

Define an elicitation method based on plurality voting as follows:

Definition: Elicitation through Plurality Voting. A set of n agents is faced with m bins, i.e., subintervals of the unit interval. Each agent chooses exactly one bin, based on their competency p_i.

Elicitation - More Competent Agents

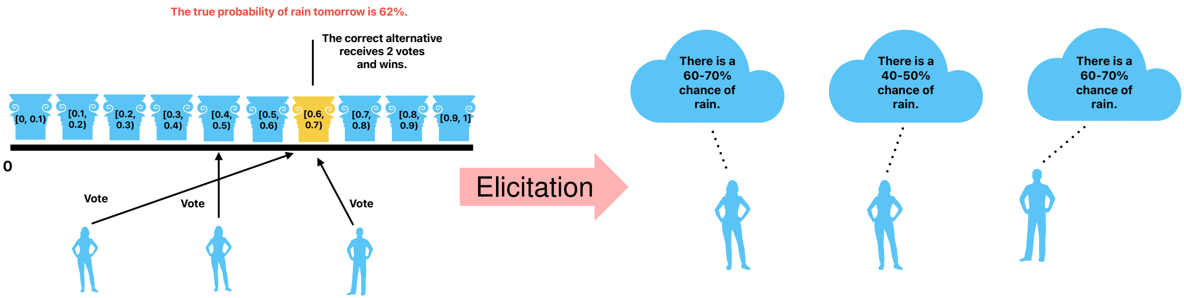

Derived a lower bound on the probability (e.g. 85%) for n independent agents (e.g. $n = 200$) choosing the correct bin over any other based on their competency p_i (e.g. $\bar{p} = 0.35$) and the number of bins m (e.g. $m = 20$).

Optimal Weights

Optimal Weights for Plurality Voting

Recall: We translated belief elicitation into a plurality voting problem.

Objective: We want to maximize the probability for the group opinion to include the correct value.

Solution: Utilize optimal weights for plurality voting[1] .

Definition (Optimal Weights.): Optimal weights for weighted plurality:

$$\lambda_i = \ln\left(\frac{(m-1)p_i^{\omega_*}}{1-p_i^{\omega_*}}\right).$$

Assuming:

- Uniform error probability: $p_i^{\omega_\dagger} = \frac{(1-p_i^{\omega_*})}{(m-1)}$,
- Competence bound: $p_i^{\omega_*} \in [\frac{1}{m}, 1]$, ensuring non-negative weights.

[1] Qing et al.: Empirical analysis of aggregation methods for collective annotation. COLING (2014).

Measure of Comparison

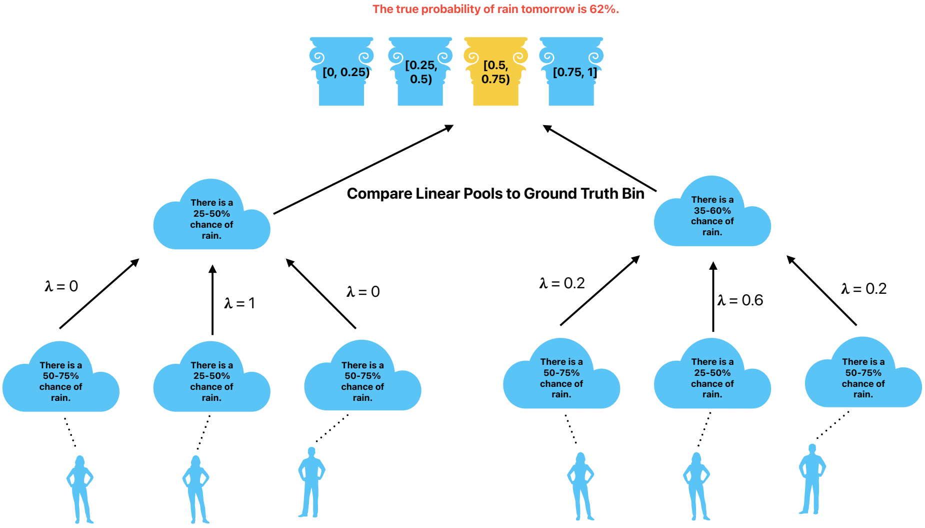

Measure of Comparison

Definition: Discrete Kullback-Leibler divergence Let $p(x)$ be the true probability distribution and $q(x)$ a model distribution for a random variable $\mathcal{X}$. The KL divergence from q to p is defined as:

$$D(p||q) = \sum_{x \in \mathcal{X}} p(x) \log \frac{p(x)}{q(x)}.$$

Example: Consider a biased coin with a 30% chance of landing heads ($p(X = 1) = 30\%, p(X = 0) = 70\%$). If an agent assumes the coin is fair ($q(X = 1) = q(X = 0) = 50\%$), the KL divergence between the true distribution and the agent's assumption is:

$$D(p||q) = p(X = 1) \log \frac{p(X = 1)}{q(X = 1)} + p(X = 0) \log \frac{p(X = 0)}{q(X = 0)} = 0.087.$$

Measure of Comparison

Definition: Imprecise Kullback-Leibler divergence Let $p(x)$ be the true imprecise probability distribution of a random variable $\mathcal{X}$, and $q(x)$ the model distribution. The Imprecise Kullback-Leibler is defined as

$$\mathcal{D}(p||q) = \frac{D(\underline{p}||\underline{q}) + D(\overline{p}||\overline{q})}{2}.$$

Side note:In imprecise probability theory, an agent's belief in proposition A is given by an interval $\mathcal{P}(A) = [a, b]$, and for its complement $\neg A$, it is $\mathcal{P}(\neg A) = [1 - b, 1 - a]$.

Example: Let $[0.2, 0.3)$ represent the aggregate obtained from linear pooling, and $[0.6, 0.7)$ represent the ground truth bin. From this, we obtain: $\mathcal{D}(\underline{p}||\underline{q}) = 0.404$, $D(\overline{p}||\overline{q}) = 0.316$, and $\mathcal{D}(p||q) = 0.36$.

Simulations

We performed experiments comparing different weight distributions across multiple parameter settings.

Four types of weights for linear pooling:

- **Uniform Weights**: identical weights across agents;
- **Bipartite Weights** [1]: Splits the agents into two competency separated groups s.t.
  - $\lambda_{lower} = \frac{1}{n} - \sigma^2 \times \frac{1}{n}$,
  - $\lambda_{upper} = \frac{1}{n} + \sigma^2 \times \frac{1}{n}$,
  - Example: $n = 200, \sigma = 0.5$, two groups of 100 agents with $\lambda_{lower} = 0.00375$, $\lambda_{upper} = 0.00625$;
- **Random weights**: generated from a uniform distribution over [0,1] and normalized;
- **Optimal weights for plurality voting**.

[1] Kriegler et al.: Imprecise probability assessment of tipping points in the climate system. PNAS 2009.

Simulations

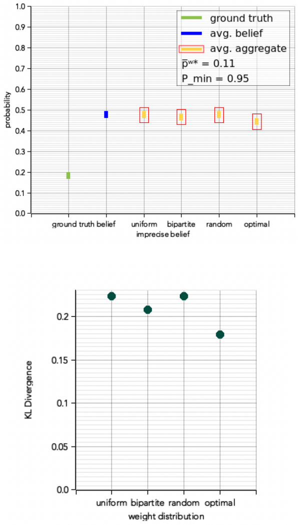

(c) $n = 2000, \bar{p}^{\omega *} = 0.11, m = 30$

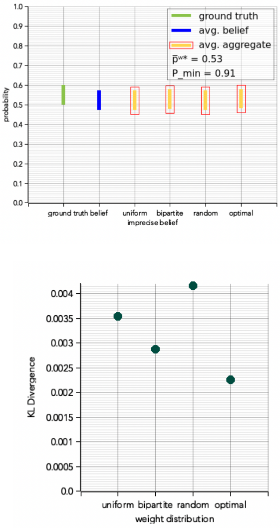

(f) $n = 40, \bar{p}^{\omega *} = 0.53, m = 10$

Summary and Next Steps

Summary:

- Translated elicitation into a plurality voting Problem;
- Derived probabilistic guarantees on the agent's beliefs quality;
- Applied optimal weights from plurality voting, and compared against weights from the literature.

Next Steps:

- Proof optimality of weights mathematically,
- Derive optimal weights for different pooling rules.