A PAC Framework for Aggregating Agents' Judgments

Hanrui Zhang Vincent Conitzer

Duke University

Collective Decision Making at Scale

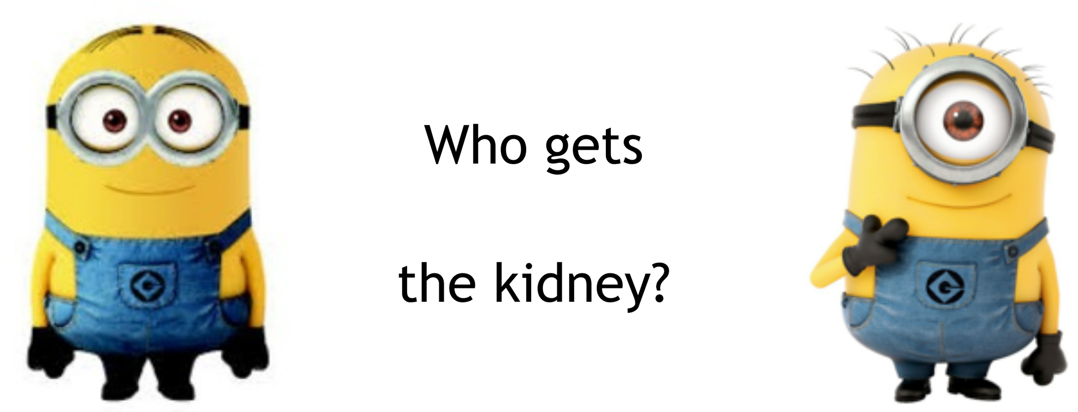

Who gets the kidney?

- 15 years old
- serious pneumonia
- doesn't drink

- 45 years old
- no other disease
- drinks every day

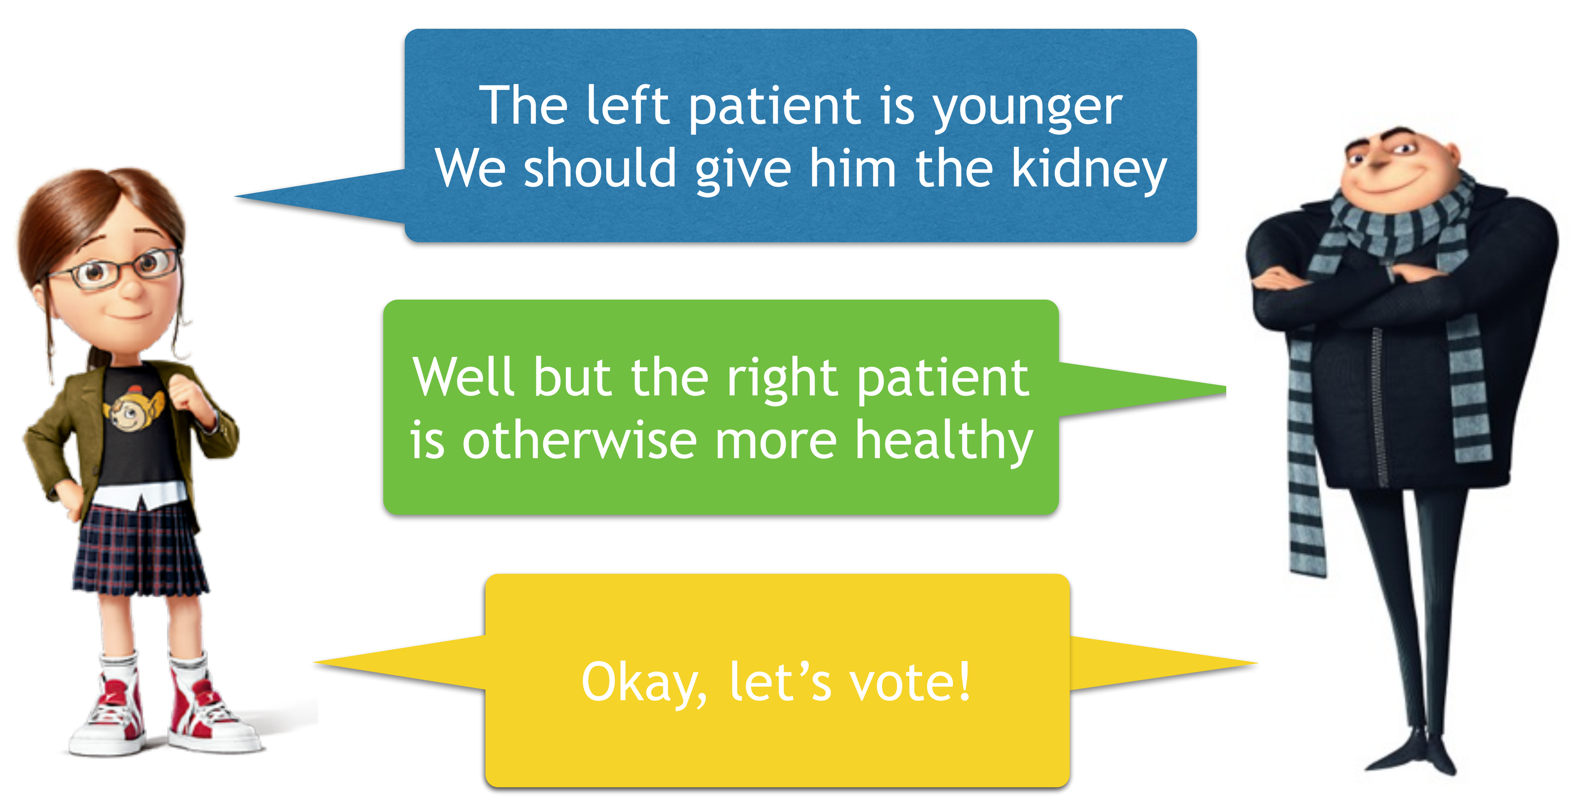

😕 But what if there are thousands of decisions to make?

💡 Maybe learn how people vote, and vote for them?

Considered in, e.g., [Noothigattu et al., AAAI'18]

Questions are:

1. How many agents do we need to query?
2. How many queries do we need to ask each of them?
3. Is there a tradeoff between the two quantities?

We study these questions from a theoretical point of view!

The PAC Aggregation Model

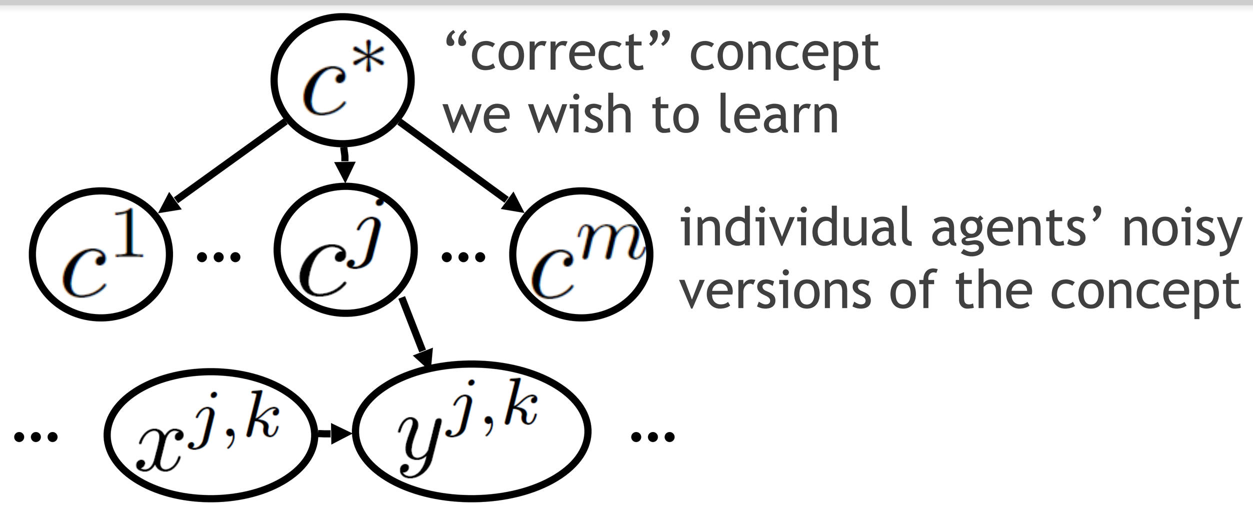

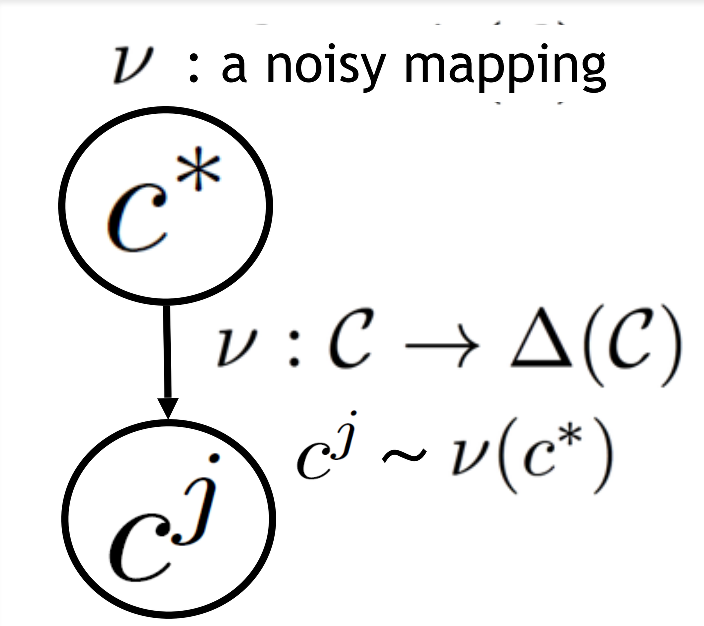

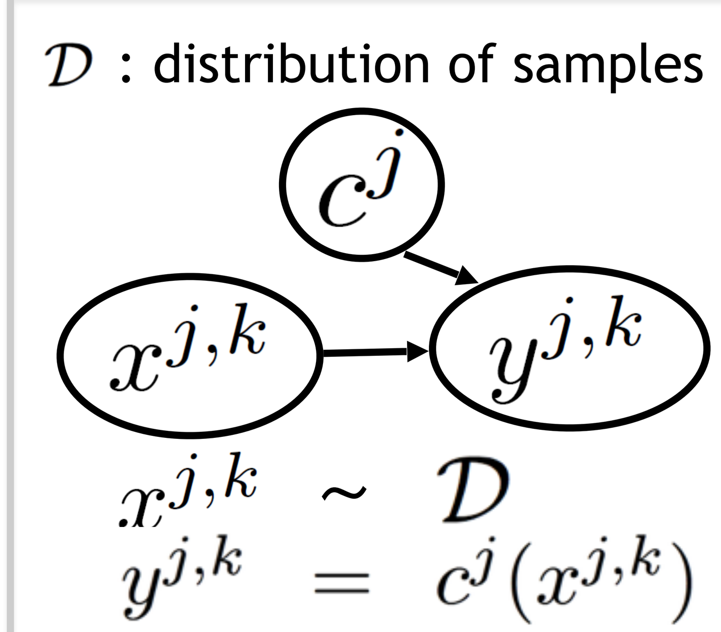

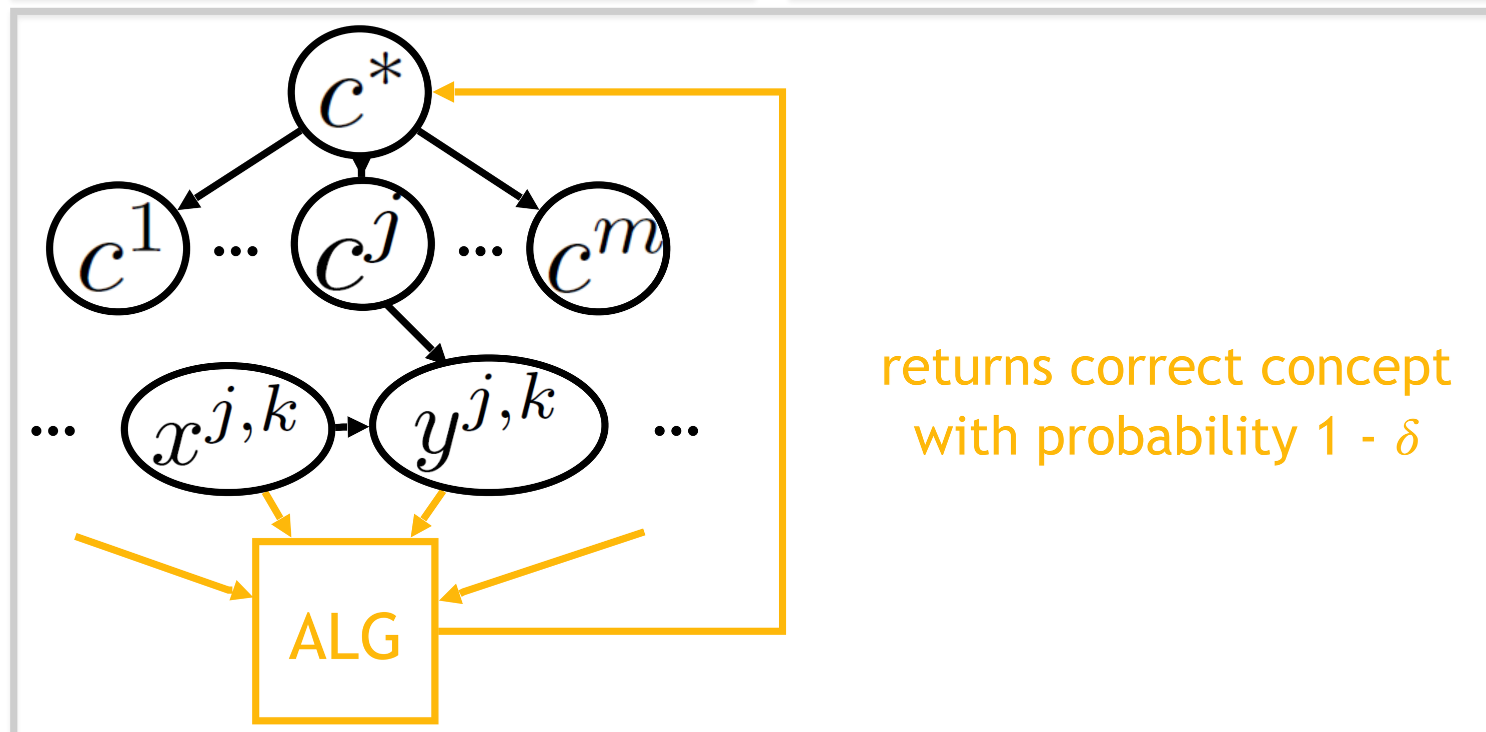

Occam's Razor for PAC Aggregation

Theorem: when (1) the minimum gap between concepts is α_{OR}, and (2) the radius of noisy mapping is β, any ERM recovers the ground truth with (1) $O(\log(|C| / \delta)\ \beta^2 / \alpha_{OR}^2)$ agents and (2) $O(\log(|C| / \delta) / \alpha_{OR}^2)$ data points, with probability $1 - \delta$

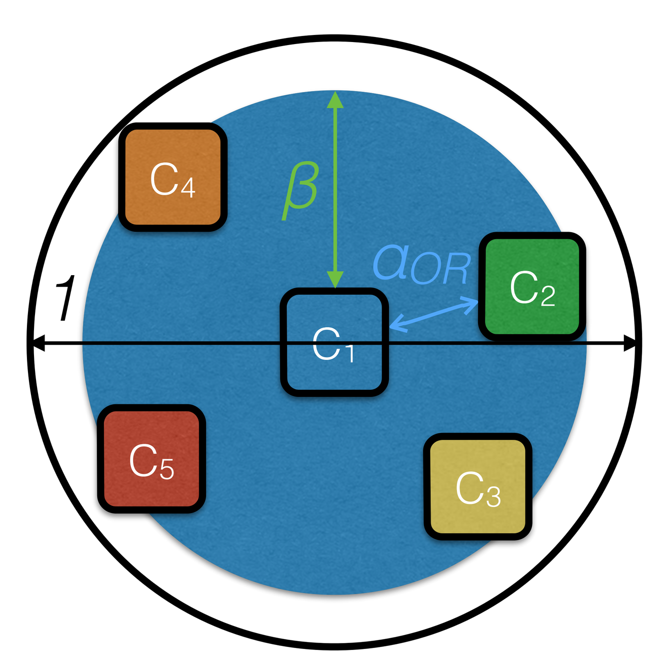

Majority Voting for PAC Aggregation

Theorem: when (1) the minimum gap between concepts is α_{MV}, and (2) the noisy version of a concept agrees with that concept with probability at least $0.5 + \theta$, the majority vote of the ERMs for each agent recovers the ground truth w.h.p.

"If the noise preserves some information, then the majority vote recovers the ground truth."

Tighter Bounds in Restricted Settings

Consider a linear model, where:

c^*, c^j, $x^{j,k}$ are n-dimensional vectors, and $y^{j,k} = \mathrm{sgn}(c^j \cdot x^{j,k})$

We give **efficient algorithms** and **matching lower bounds,** when, e.g.:

(1) each coordinate of a concept is +/-1, (2) the distribution is symmetric, and (3) the noise flips each coordinate w.p. 0.5 η

Theorem: there is an algorithm that recovers the ground truth with $O(\log(n / \delta) / (1 - \eta)^2)$ agents and $O(n \log(n / \delta) / (1 - \eta)^2)$ data points, with probability $1 - \delta$, and this is tight up to a constant / $\log(n / \delta)$ factor.

Paradigms of Aggregation