Epistemic Sets Applied to Best-of-n Problems

Jonathan Lawry

Department of Engineering Mathematics, University of Bristol, UK

ECSQARU 2019

Talk Outline

- Distributed learning in multi-agent systems: The two processes of evidential updating and fusion.
- The Best-of-n problem.
- An agent-based discrete time model with epistemic sets.
- Fusion of epistemic sets and negative updating.
- Simulation experiments exploring robustness to noise and scalability.
- Mathematical model of the system based on ODEs.
- Conclusions

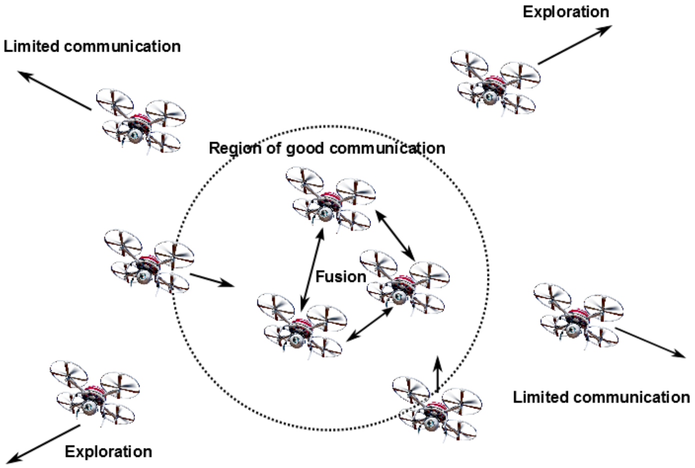
Exploration
Limited communication
Region of good communication
Fusion
Limited communication
Exploration

Two Processes: Fusion and Evidential Updating

- We consider distributed learning from the perspective of two interacting processes:
- Direct learning from the environment.
- Learning from other agents.

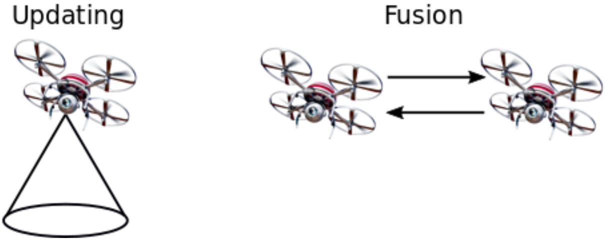

The Best-of-n Problem

- The best-of-n problem is a broad class of distributed learning problem where the aim is for a population of agents to identify which out of *n* options is the best.
- This is achieved by direct feedback from the environment and agents sharing information.

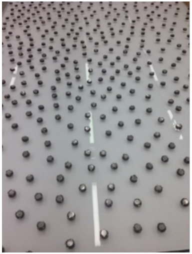

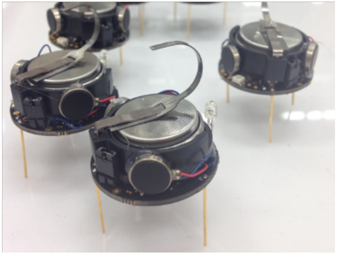

A Simple Agent-Based Model

- Let $\mathbb{S} = \{s_1, \ldots, s_n\}$ the set of all possible states of the world.
- For s_i there is an associated quality value $q_i \in [0, 1]$.
- Without loss of generality we will assume that the states are enumerated so that $q_1 < q_2 < \ldots < q_n$.
- There is a population of k agents where each agent's belief is represented by an epistemic set $B \subseteq \mathbb{S}$ such that $B \neq \emptyset$.
- Agents are initialized with beliefs $B = \mathbb{S}$ representing complete ignorance.
- We adopt a discrete time model where at each time step two agents are selected at random to fuse their beliefs.
- Also, at each time step each agent chooses to investigate two particular states.
- This investigation will provide evidence with probability ρ (the evidence rate).

Fusion of Epistemic Sets

Definition

Intersection & Union Pooling Operator
For $\emptyset \neq B_1, B_2 \subseteq \mathbb{S}$*;*

$$B_1 \odot B_2 = \begin{cases} B_1 \cap B_2 : B_1 \cap B_2 \neq \emptyset \\ B_1 \cup B_2 : B_1 \cap B_2 = \emptyset \end{cases}$$

- This is the only operator for combining epistemic sets which satisfies the following properties:
- *Optimism*: If $B_1 \cap B_2 \neq \emptyset$ then $B_1 \odot B_2 \subseteq B_1(B_2)$.
- *Unanimity*: $B_1 \cap B_2 \subseteq B_1 \odot B_2 \subseteq B_1 \cup B_2$.
- *Minimal Commitment*: $B_1 \odot B_2$ is the largest epistemic set satisfying both *optimism* and *unanimity*.

Negative Updating

- We assume that evidence is received in the form of a direct comparison between the quality values of two states.
- E is a comparison between the quality values for states s_i and s_j and if $q_i > q_j$ then this can be represented by the epistemic set $E_j = \mathbb{S} - \{s_j\}$ expressing the information that s_j is not the best state.

Definition

Evidential Updating
For $\emptyset \neq B, E \subseteq \mathbb{S}$;

$$B|E = \begin{cases} B \cap E : B \cap E \neq \emptyset \\ B : B \cap E = \emptyset \end{cases}$$

Updating and Noise

- At each time step every agent picks a two distinct states at random from B to investigate.
- If $|B| = 1$ no investigation is carried out.
- With probability $1 - \rho$ no evidence is found and the agent's belief remains unchanged.
- If states s_i and s_j are sampled and $q_i > q_j$ then if qualities are sampled correctly the agent should update to $B|E_j$.
- We assume that there is a probability of sampling error taken to be a decreasing function of $q_i - q_j$ so that:

$$B|E = \begin{cases} B|E_j : \text{with probability } 1 - f(q_i - q_j) \\ B|E_i : \text{with probability } f(q_i - q_j) \end{cases}$$

- We assume that quality values are uniformly distributed over the interval $[0, 1]$ so that $q_i = \frac{i}{n+1}$ for $i = 1, \ldots, n$.

An Error Model

For $d = q_i - q_j$;

$$f(d) = \begin{cases} \frac{0.5\left(e^{-\lambda d} - e^{-\lambda}\right)}{1 - e^{-\lambda}} : \lambda \neq 0 \\ 0.5(1-d) : \lambda = 0 \end{cases}$$

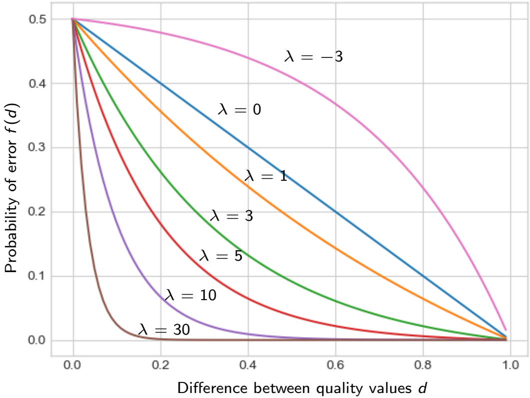

It's Good to Talk

- Does fusion aid learning in the no noise case?

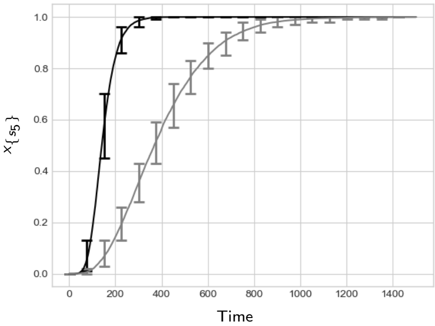

- Time series of the proportion of agents with belief $\{s_5\}$ averaged over 100 runs of the simulation. Parameter values are $k = 100$, $n = 5$ and $\rho = 0.01$.
- The grey line is for evidential updating only and the black line is for both updating and pooling.

Fusion and Robustness

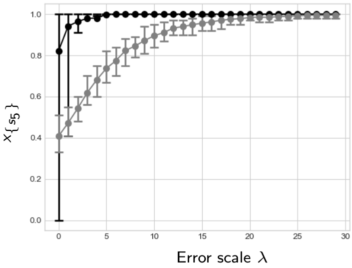

(a) Figure A

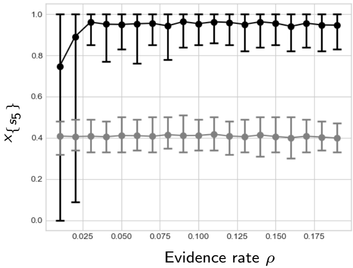

(b) Figure B

- 100 simulation runs with $k = 100$ and $n = 5$. All results are after 1500 iterations.
- Figure A: Average $x_{\{s_5\}}$ against λ. Here $\rho = 0.01$.
- Figure B: Average $x_{\{s_5\}}$ against ρ. Here $\lambda = 0$

Robustness and Scalability

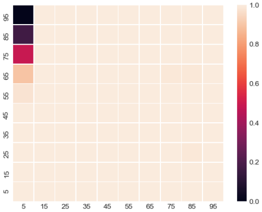

(c) No Noise

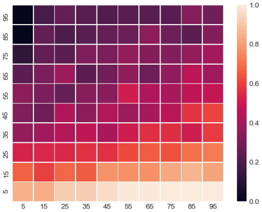

(d) $\lambda = 10$

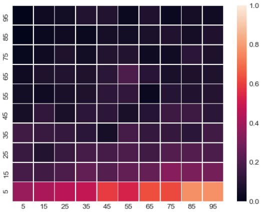

(e) $\lambda = 0$

Robustness and Scalability II

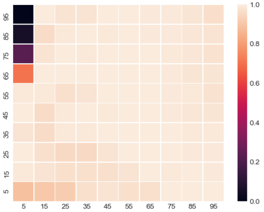

(f) $\lambda = 10$, $\sum_{i:q_i \geq \frac{3}{4}} x_{\{s_i\}}$

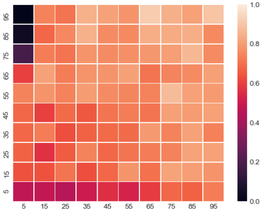

(g) $\lambda = 0$, $\sum_{i:q_i \geq \frac{3}{4}} x_{\{s_i\}}$

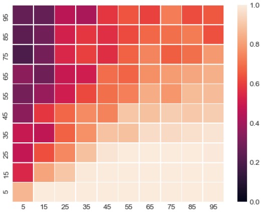

(h) $\lambda = 10$, $\rho = 0.1$

Analytical Model

- For a fixed number of states we can write down a mathematical model of this multi-agent scenario as a system of ordinary differential equations.
- For $n = 3$ this gives a system of 7 ODEs with 6 degrees of freedom since $\sum_{B \neq \emptyset} x_B = 1$.

$$\begin{aligned} \dot{x}_{\{s_1,s_2,s_3\}} &= 2\pi(1-\rho)x_{\{s_3\}}x_{\{s_1,s_2\}} + 2\pi(1-\rho)x_{\{s_2\}}x_{\{s_1,s_3\}} \\ &\quad +2\pi(1-\rho)x_{\{s_1\}}x_{\{s_2,s_4\}} \\ &\quad +x_{\{s_1,s_2,s_3\}}\left(-(1-\pi)\rho - \pi + \pi(1-\rho)x_{\{s_1,s_2,s_3\}}\right) \\ &\vdots \\ &\vdots \end{aligned}$$

- Here π is the probability that an given agent takes part in fusion. In our case $\pi = \frac{2}{k} - \frac{1}{k^2}$.

Fixed Point Analysis

- The fixed points for this system of equations are those values of $x_B : B \neq \emptyset$ satisfying $\dot{x}_B = 0$.
- Stability is then determined from the Jacobian $J = \left(\frac{\partial \dot{x}_i}{\partial x_j}\right)$ by evaluating J at each fixed point and finding the eigenvalues.
- A fixed point is stable if and only if the real parts of all the eigenvalues are negative.

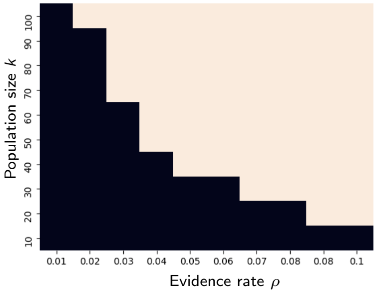

Model Based Simulations

- We simulation this system of equation such that $x_B^{t+1} = x_B^t + \Delta t \dot{x}_B$.

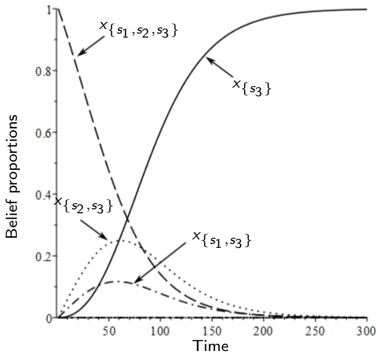

- Trajectories of the proportions holding different beliefs as function of time generated from the ODEs with $k = 100$, $\rho = 0.01$ and initialised so that $x_{\{s_1,s_2,s_3\}} = 1$ at time $t = 0$.

Conclusions

- Epistemic sets are a very simple way of representing uncertain beliefs in AI which provide a computationally efficient approach to belief updating and fusion.
- Despite their representational limitations we have shown that they can provide a framework in which a whole population of agents can efficiently solve best-of-n problems.
- By combining updating based on direct evidence with belief pooling between agents, the agent population is able to compensate for sparsity of evidence, and also correct errors resulting from noise in the evidence collection process.
- In future work we will investigate the application of epistemic sets to a broader class of decentralised learning problems, operating under different local interaction rules.