Bayes Bots:

Bayesian Decision-Making for Robot Swarms

Julia Ebert,[1] Melvin Gauci,[1,2] and Radhika Nagpal[1,2]
[1] Harvard University, [2] Wyss Institute

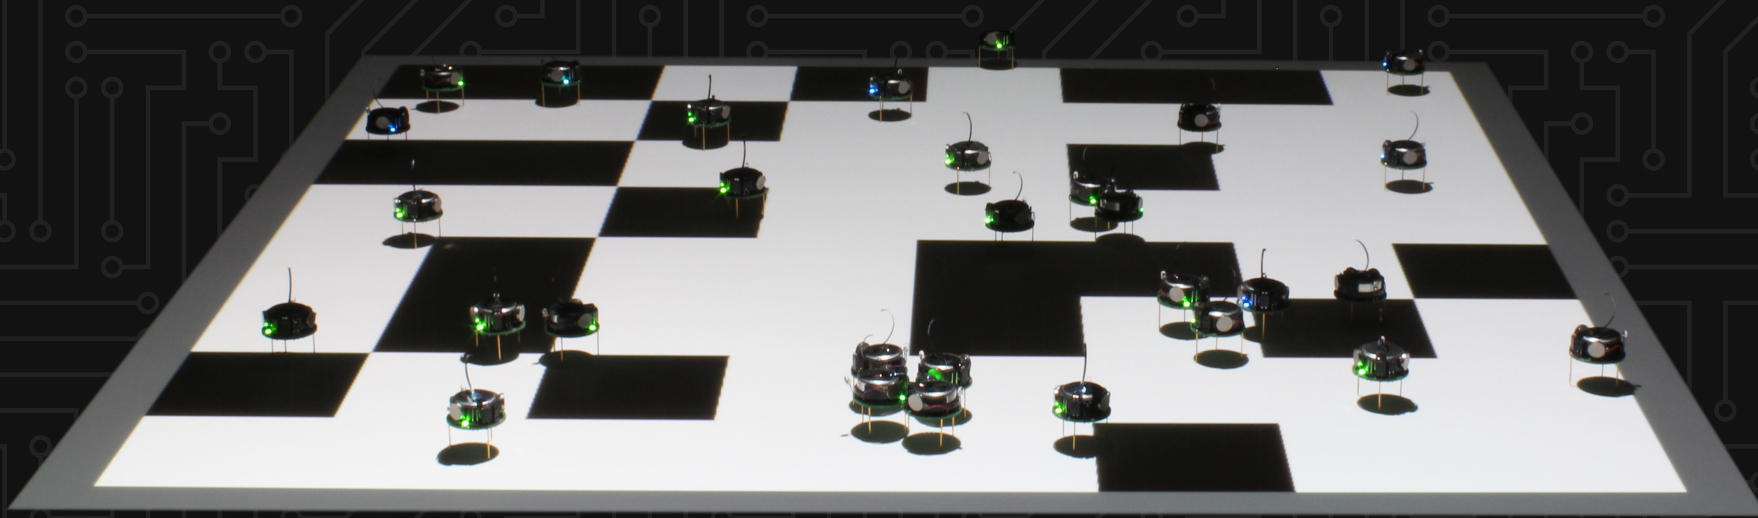

Background

GENERAL PROBLEM
How can a swarm of robots collectively make accurate and fast decisions about features of their environment?

CHALLENGE
Robots only have local observations and communication

CASE STUDY
Kilobots deciding whether an environment is mostly black or mostly white

SOLUTION
Robots maintain Bayesian model of environment, update with observations, and apply decision rule.

Bayesian decision-making provides a leaderless, mathematically-grounded decision framework that can be applied across robot and environment complexities.

Agent Model & Algorithm

BAYESIAN MODEL
Prior distribution of fill ratio f: $f \sim \text{Beta}(\alpha, \beta)$
Likelihood of color C: $C \sim \text{Bernoulli}(f)$
Posterior after observing color: $f \mid C \sim \text{Beta}(\alpha + C, \beta + (1 - C))$

1. MOVEMENT
Continuous pseudo-random walk in bounded arena

2. OBSERVATION
Observe black/white color C after observation interval

3. POSTERIOR UPDATE
Update posterior with own and received observations

4. DECISION
Commit when sufficient credible interval (credible threshold) of posterior is above or below 0.5

5. COMMUNICATION
- *Transmit* most recent observation OR decision
- *Receive* observations from other robots and update posterior

Simulations

SETUP
100 Kilobot robots in 2.4 m x 2.4 m arena in the Kilosim simulator
100 trials per condition (5,280 parameter combinations)

PARAMETER SWEEPS
Symmetric Beta prior: { 1, 2, 3, 5, 8, 10, 12, 15, 25, 30 }
Positive feedback: { 0, 1 }
Observation interval (s): { 1, 5, 10, 15, 30, 45, 60, 90, 120, 150, 200 }
Credible threshold: { 0.9, 0.95, 0.98, 0.99 }
Fill ratio: { 0.52, 0.55, 0.6, 0.7, 0.8 }

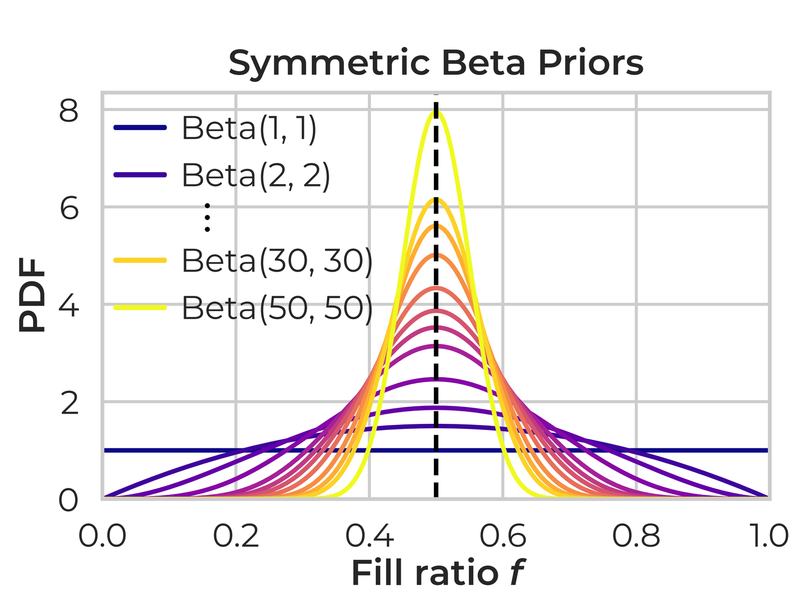

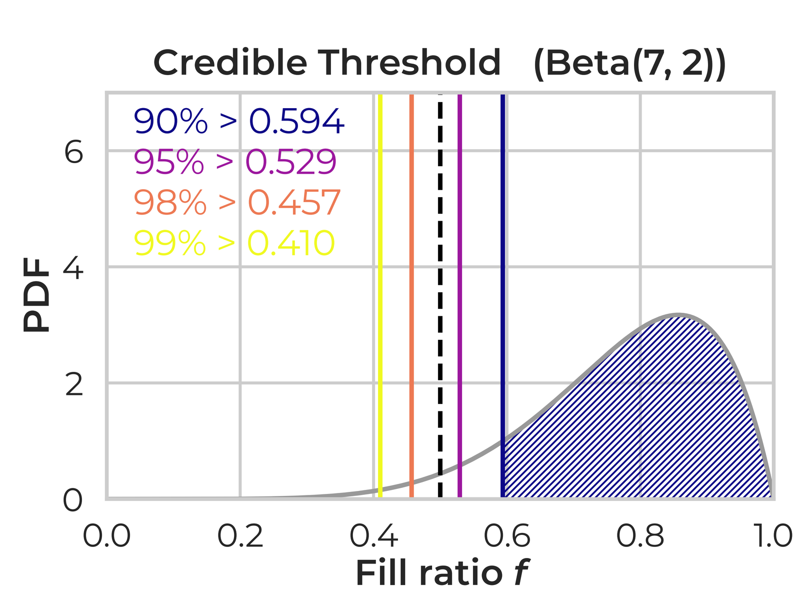

Results

GOAL
Understand the **speed/accuracy trade-off** in decision-making as a multi-objective optimization problem by comparing against the Pareto front for a fill ratio of 0.52.

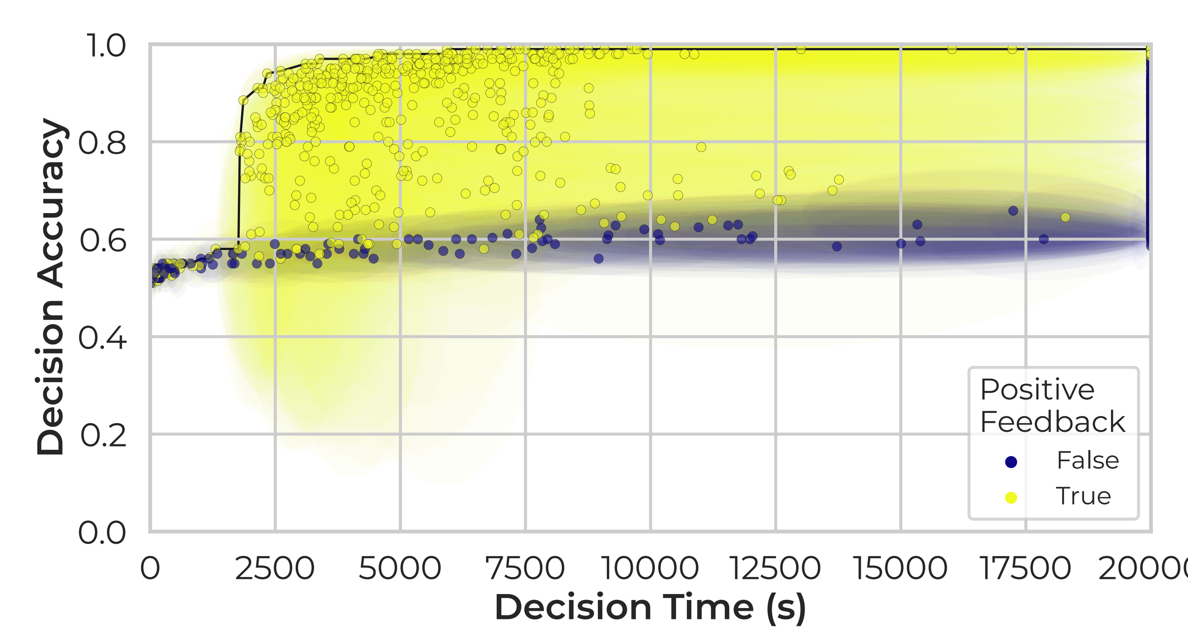

FEEDBACK
Effect: Using bio-inspired positive feedback results in dramatically faster and higher accuracy decisions

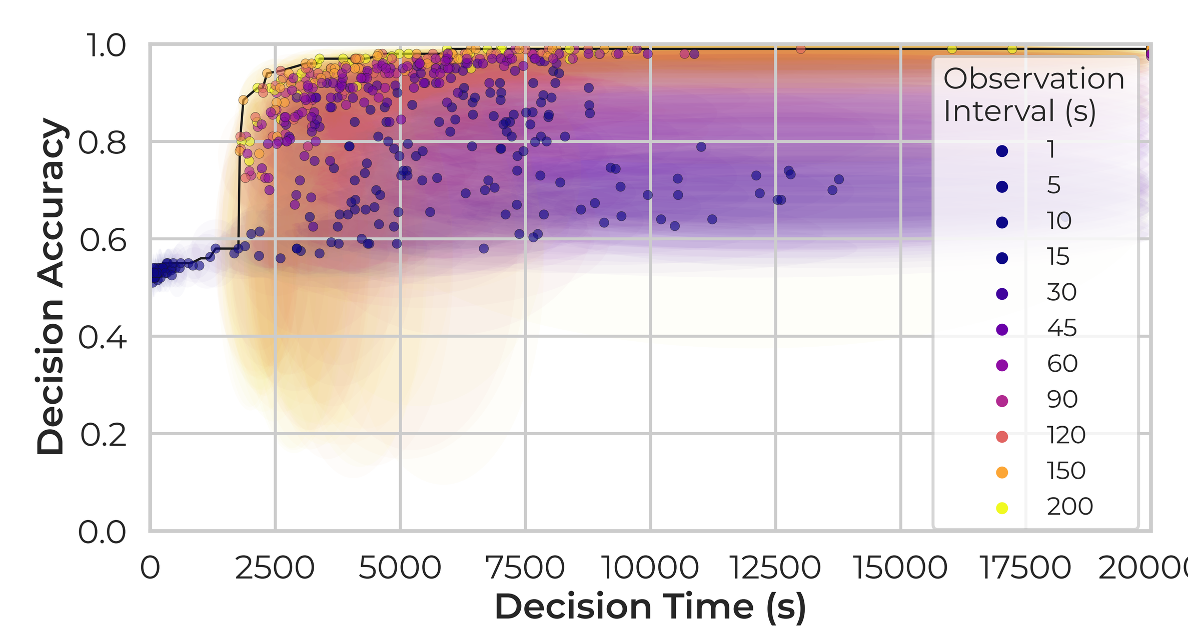

OBSERVATION INTERVAL
Effect: Surprisingly, longer times between observations are closer to the Pareto front; increased spatial mixing decreases the total swarm decision time
Showing only: Conditions with positive feedback

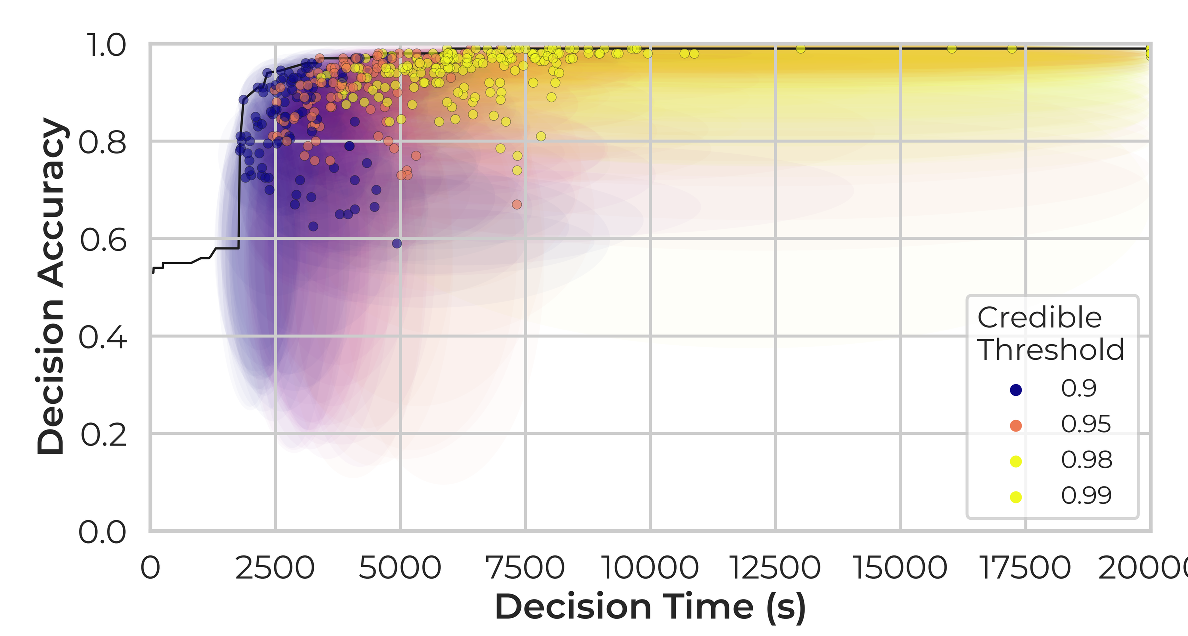

CREDIBLE THRESHOLD
Effect: Lower credible threshold saves time with minimal accuracy cost
Showing only: Conditions with observation interval ≥ 15 s

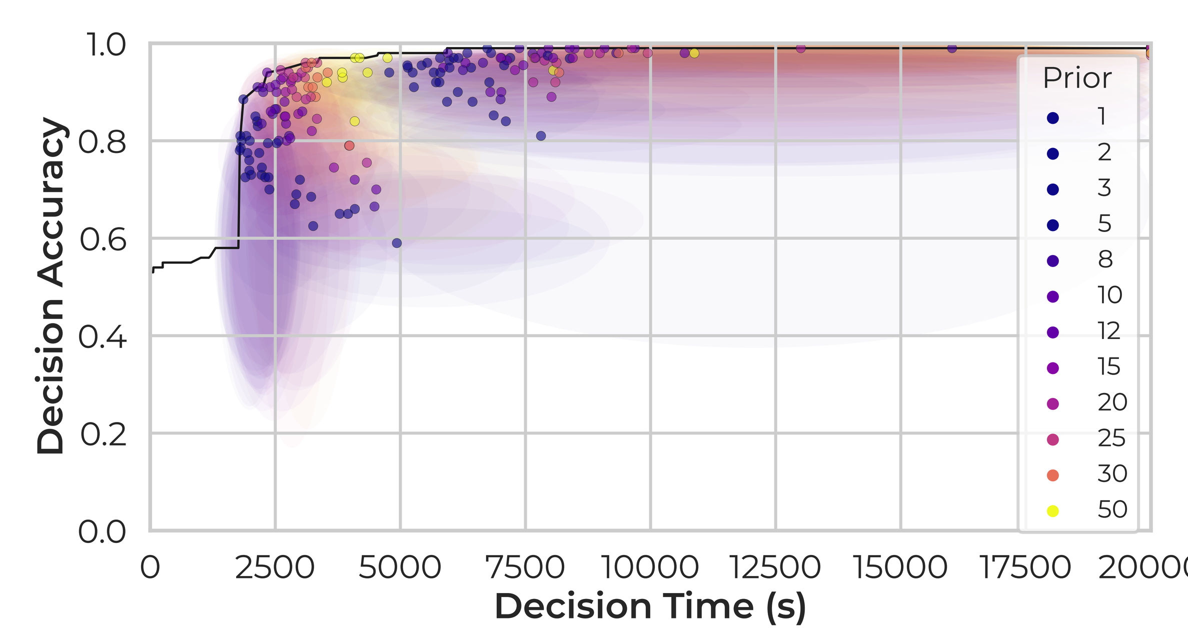

PRIOR
Effect: Lower credible thresholds are effective only if a regularizing prior prevents premature decisions
Showing only: Conditions with credible threshold of 0.9 and 0.99

Discussion & Future Work

SUMMARY
We show a "cheap lunch" effect with tunable (but non-intuitive) trade-offs:
- Positive feedback improves decision-making, rather than creating bifurcations
- Robots making *fewer* observations improves accuracy by reducing spatial effects
- Selecting a sufficient regularizing prior allow for a lower credible threshold with a small time cost

FUTURE WORK
- Extend positive feedback mechanism to more complex *informed communication*
- Add informed movement (adaptive sampling) instead of random walks
- Extend to multiple features with multi-dimensional distributions (e.g., Dirichlet)
- Generalization to more complex robots and environmental features
- Compare to previous bio-inspired decision-making algorithms and ongoing theoretical work

REFERENCES
JT Ebert, M Gauci, & R Nagpal (2018). "Multi-Feature Collective Decision Making in Robot Swarms." *AAMAS*.

FUNDING
Department of Energy Computational Science Graduate Fellowship (Ebert), Wyss Institute Technology Development Fellowship (Gauci)

ACKNOWLEDGEMENTS
Richard Barnes, for parallelization and optimization on Kilosim; Frederik Mallman-Trenn, for theory work on provable collective decision-making