Infotaxis in a Multi-agent Sensor Network

Julia Ebert[1], Clark Teeple[2], Emma Steinhardt[2], and Sharad Ramanathan[3]
[1]Computer Science, [2]Mechanical Engineering, [3]Applied Physics, Harvard University

HARVARD John A. Paulson School of Engineering and Applied Sciences

BACKGROUND

PROBLEM

How do we find a source when observations are sparse?

CHALLENGES:

Observations are too sparse for gradient descent

SOLUTION:

Infotaxis: Follow the information gradient

SETUP

ENVIRONMENT:

Particles emitted from source at constant rate

Confined, discretized arena with constant wind velocity

AGENTS:

Goal: reach source

Particle detection ability

Move one cell per tick

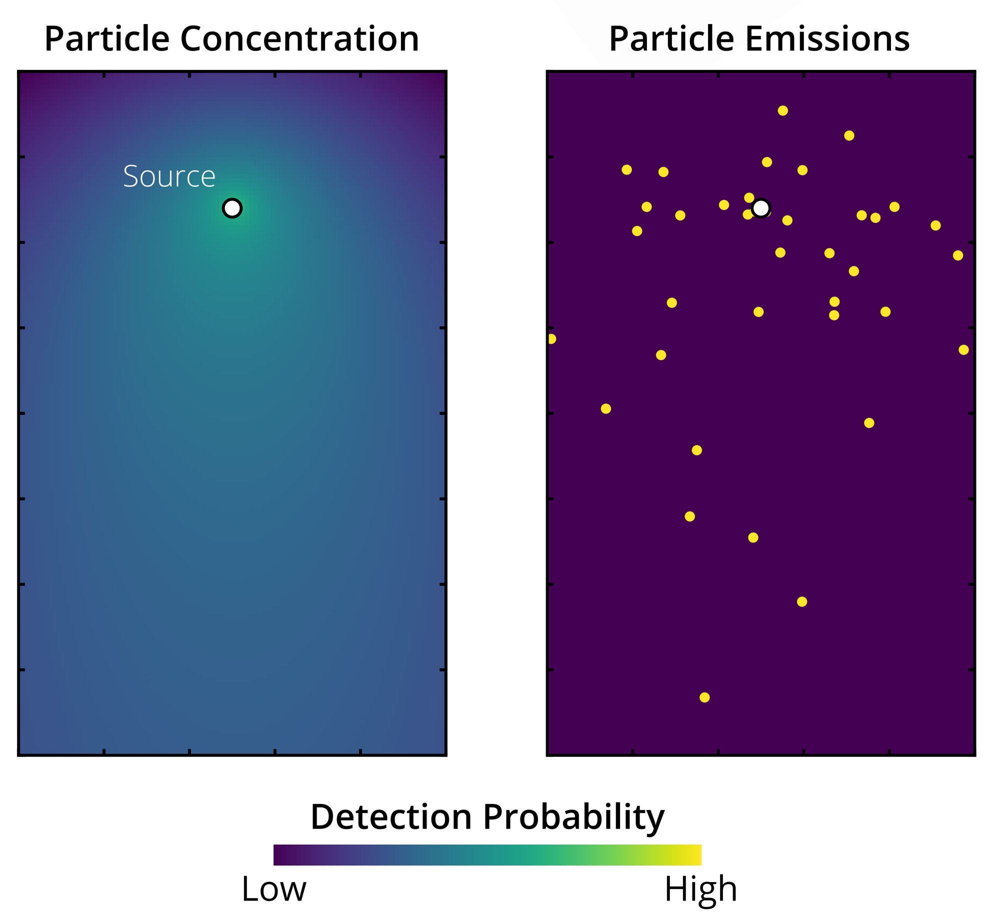

SINGLE-AGENT INFOTAXIS

PROBABILITY MAP

Agents maintain map of the likelihood that the source is in each cell

Observation history (trace):

$$\Gamma_t = \{(r_a, t_a)\}$$

Likelihood that source is at $\mathbf{r}_0$:

$$R(\mathbf{r}|\mathbf{r}_0) = \text{estimated particle detection rate at } \mathbf{r} \text{ if source is } \mathbf{r}_0$$

$$L(\Gamma_t|\mathbf{r}_0) = \underbrace{\exp\left(\sum_a \sum_{t' \in T_a} R(\mathbf{r}_a(t')|\mathbf{r}_0)dt'\right)}_{\text{Prior}} \prod_{b=1}^{T_b} R(\mathbf{r}(t_b), t_b|\mathbf{r}_0)$$

Probability map over full arena:

$$P_t(\mathbf{r}_0|\Gamma_t) = \frac{L(\Gamma_t|\mathbf{r}_0)}{\sum_q L(\Gamma_t|\mathbf{r}_q)d\mathbf{r}_q}$$

MOVEMENT DECISION

Agents move to cell that maximizes entropy

Expected information gain:

$$\Delta S = \text{Information from expected detections at } \mathbf{r}_j$$
$$\text{(Poisson distribution)}$$

$$H = \text{Entropy of } P_t(\mathbf{r}_j)$$

$$\Delta H(\mathbf{r}_i \to \mathbf{r}_j) = \underbrace{P_t(\mathbf{r}_j)(-H)}_{\text{Exploitation}} + \underbrace{[1 - P_t(\mathbf{r}_j)]\,\Delta S}_{\text{Exploration}}$$

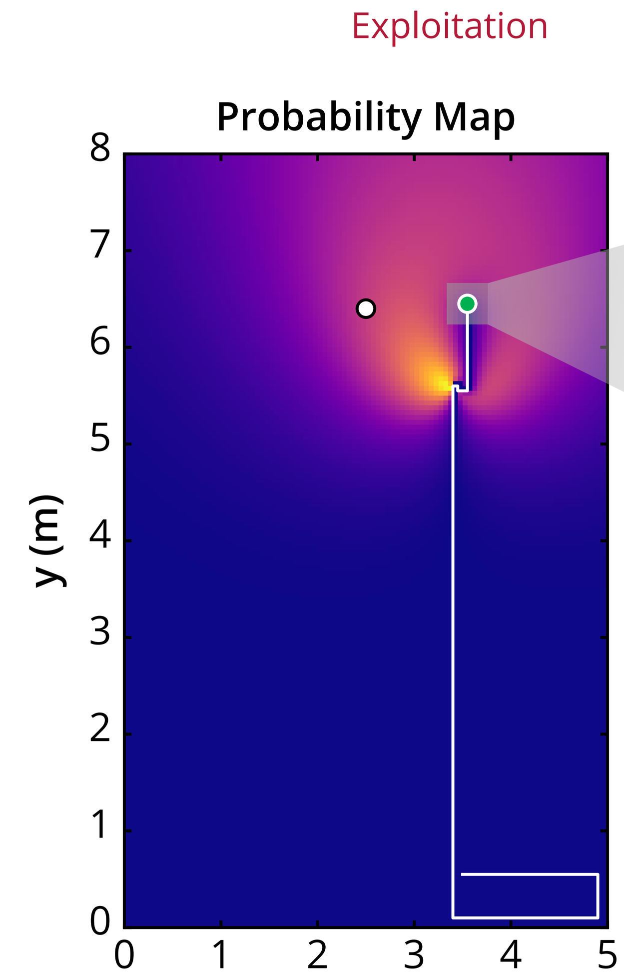

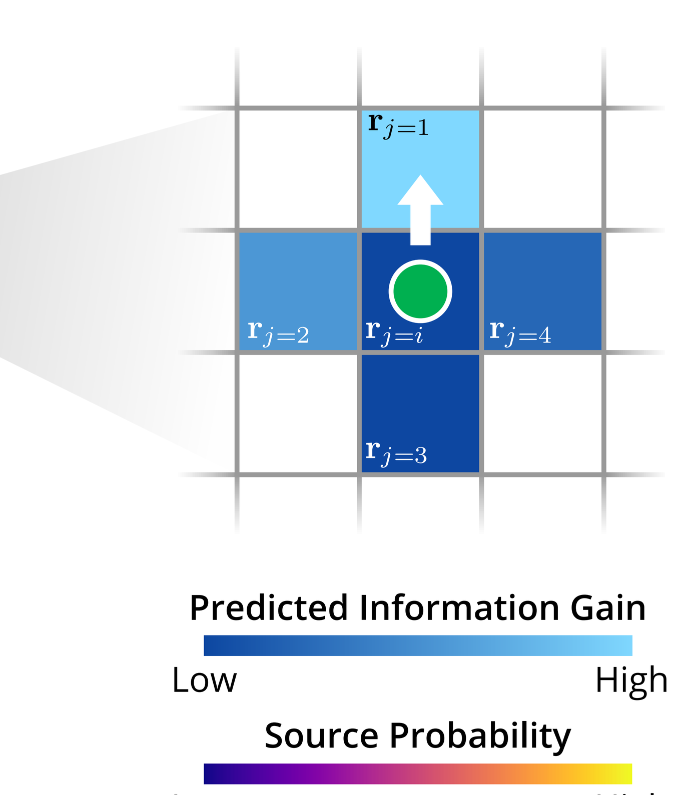

Vergassola, Villermaux, & Shraiman (2007). "Infotaxis as a strategy for searching without gradients." *Nature.*

MULTI-AGENT INFOTAXIS

COLLISION AVOIDANCE

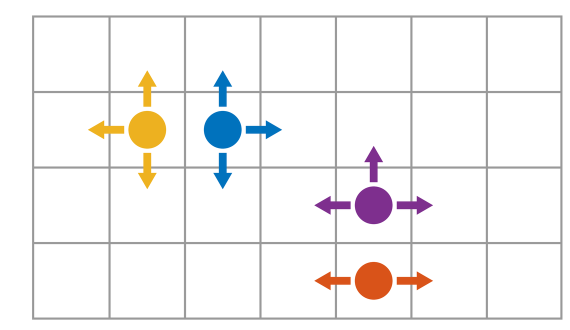

COMMUNICATION

Continuous communication over fixed range

Agents treat others' observations as their own

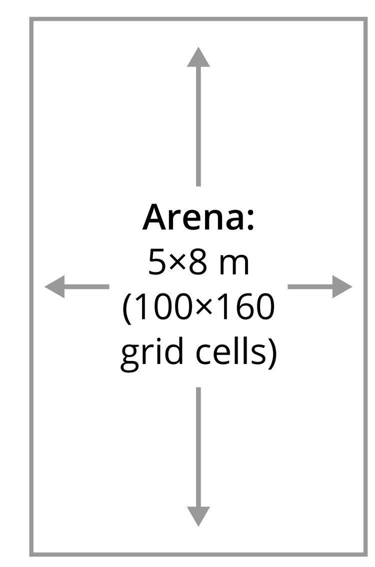

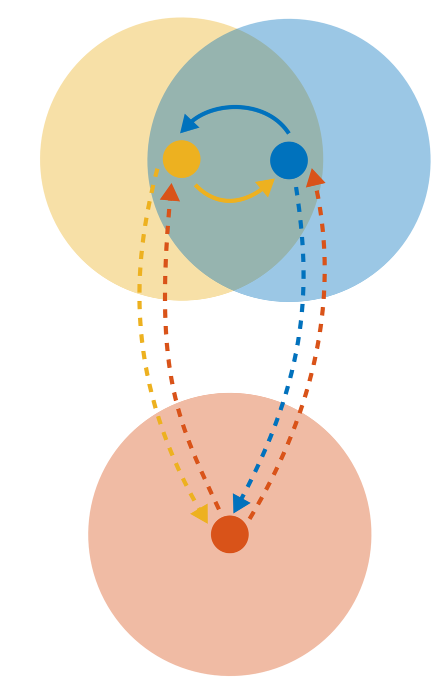

EXPERIMENTS

VARY:

Communication range

Number of agents

INVESTIGATE:

Time for first agent to reach source

RESULTS

AGENT TRAJECTORIES

Sweeping search patterns as observed in moths and dogs

Early Behavior: Exploration dominates; agents spread out

Late Behavior: Exploitation dominates; agents cluster

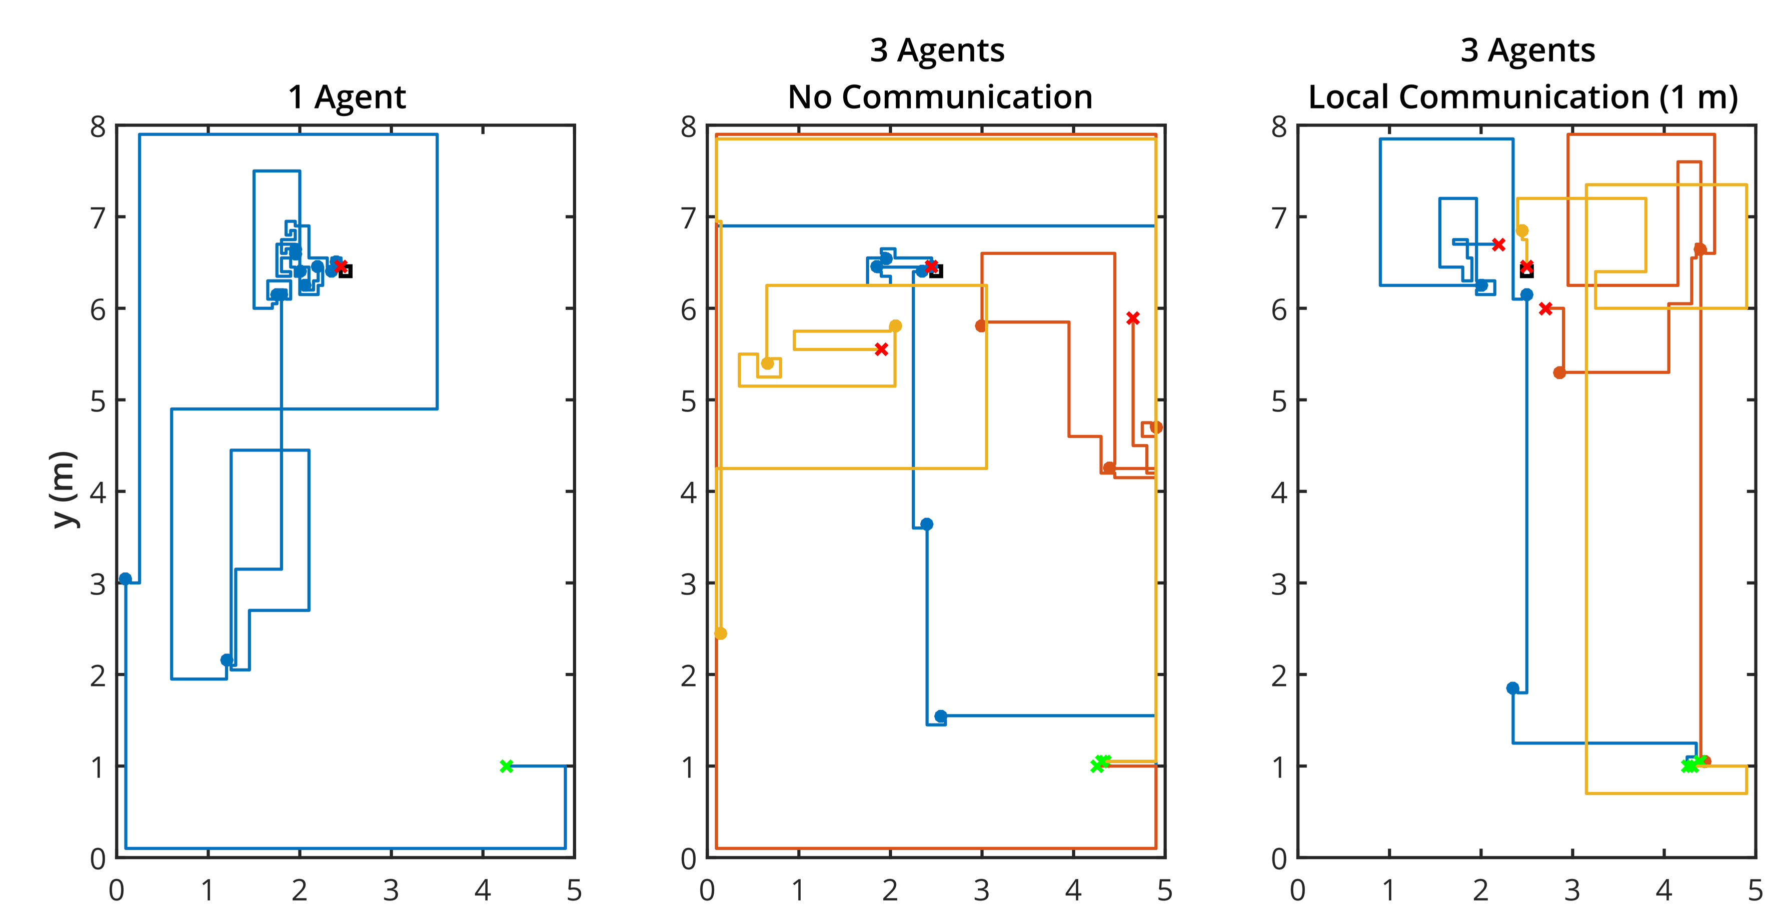

MULTI-AGENT EFFECTS

Communication Range:
Gains for local communication

Number of Agents:
Diminishing returns for increasing number of agents

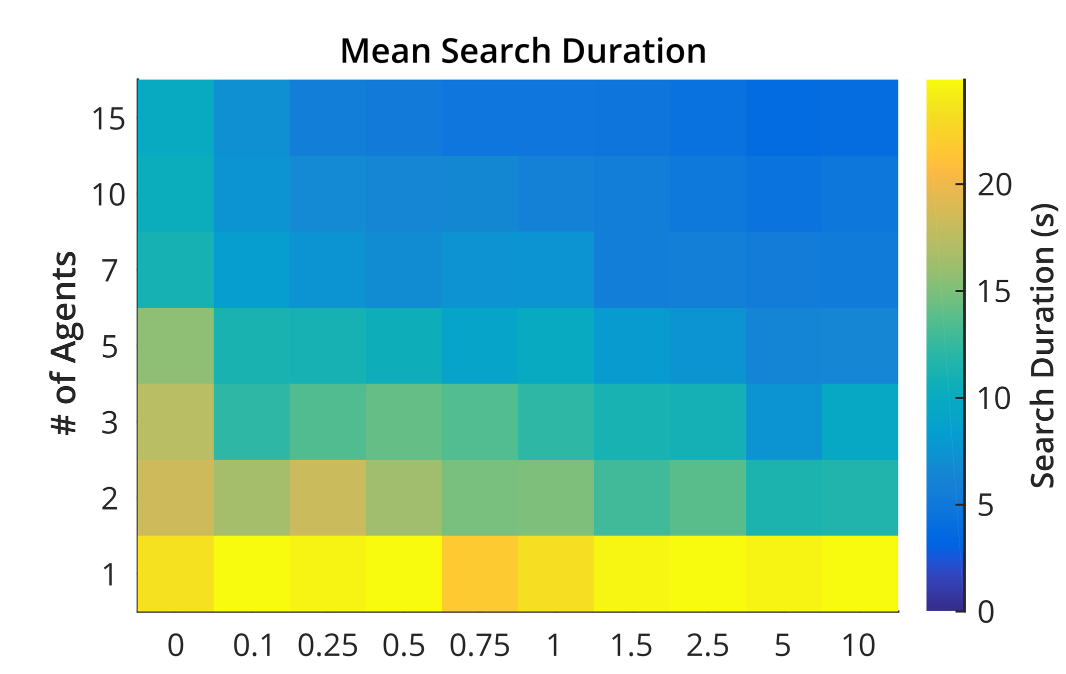

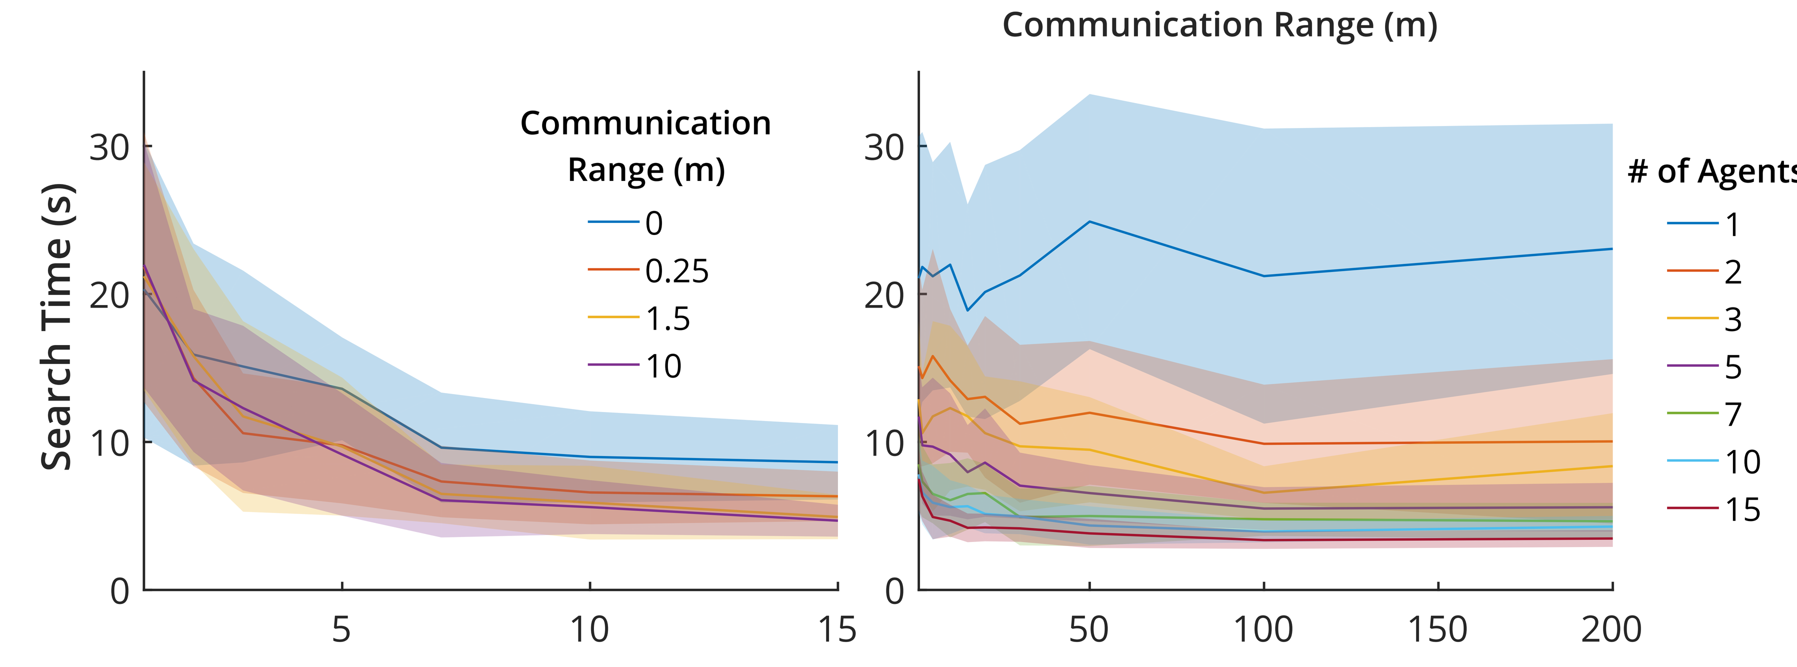

FUTURE WORK

Multi-agent behavior: Quantify group dynamics, test robustness

Reduce assumptions: Lossless communication, perfect localization

Localization: Noisy, non-discrete, filtered

Communication: Asynchronous, lossy

Computational complexity: Continuous probability map (mixture model)