Optimal Bayesian olfactory search in a realistic turbulent flow

R. A. Heinonen[1], F. Bonaccorso[1], L. Biferale[1], A. Celani[2], and M. Vergassola[3]

[1]Dept. Physics and INFN, University of Rome, "Tor Vergata"
[2]The Abdus Salam International Center for Theoretical Physics
[3]Dept. Physics, Ecole Normale Supérieure

Supported by the European Research Council under grant No. 882340

European Research Council
Established by the European Commission
Supporting top researchers from anywhere in the world

Introduction

Introduction: searching for an odor source in a turbulent environment

- Problem: find source of odor or other cue advected by atmosphere (e.g. moth drawn to mate by pheromones)
- Turbulence mixes cue into intermittent landscape: randomizes cue encounters, mean conc. gradients slow to converge
- How to search without gradients?

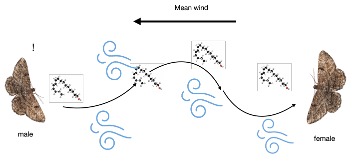

Figure Artist's conception of a moth searching for a mate via pheromone cues.

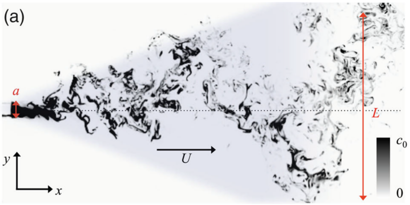

Figure Concentration field from jet flow experiment [Villermaux and Innocenti, 1999]. Fig taken from [Celani et al., 2014]

Model search problem

- Source fixed on 2D grid, location unknown to agent.
- At each Δt, agent makes observation $o(t) = \theta(c(t) - c_{\mathrm{thr}})$ then moves. Start with a detection at $t = 0$
- Try to reach source in as few Δt as possible
- Key physics input is $\Pr(o|\mathbf{s})$, where $\mathbf{s} \equiv \mathbf{r} - \mathbf{r}_0$.

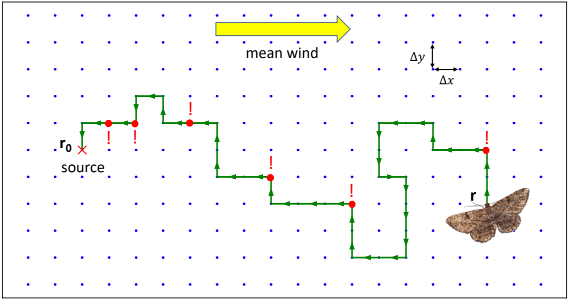

Figure In our setup, agent lives on the gridworld (blue points) and tries to find the source (red x)

Bayesian approach

- Assuming independent observations, agent can store information in a probability distribution (posterior) $b(\mathbf{s})$ over possible source locations
- After each observation, update posterior using Bayes' theorem

$$b(\mathbf{s}'|o) \propto b(\mathbf{s})\mathrm{Pr}(o|\mathbf{s})$$

- Try to find policy $\pi : b \mapsto \mathbf{a}$ minimizing expected time of arrival to source
- N.B.: $\mathrm{Pr}(o|\mathbf{s})$ assumed known by agent

Infotaxis: an important heuristic

- [Vergassola et al., 2007] suggested a policy that seeks to maximize information content of belief

$$\pi(b) = \arg\min_a \sum_o \Pr(o|b,a) H[b_{o,a}]$$

  where $H[b] = -\sum_s b(s) \log b(s)$.
- Generally performs extremely well, but can improve by adding information about distance from source ("space-aware infotaxis") [Loisy and Eloy, 2022]

Optimal policy

- One can show that the optimal policy (minimum mean arrival time) satisfies a so-called the *Bellman* equation, which can be solved algorithmically
- Recent work solved the problem using three algorithms (Perseus w/ reward shaping, SARSOP, model-based DQN). Can usually beat all available heuristics
  1. Loisy and Eloy *Proc. R. Soc. Lond.* (2022) — DQN in windless setting
  2. RAH, Biferale, Celani, and Vergassola *PRE* (2023) — Perseus in windy setting
  3. Loisy and RAH *EPJE* (2023) — benchmark on Perseus, SARSOP, DQN in windy and windless settings
- But this work done in a 'toy model' setting with artificial detections

Correlations

- Real flows will exhibit correlations between successive observations: $\Pr(o_t|\mathbf{s}) \neq \Pr(o_t|o_{t-1}, \mathbf{s})$
- Correlations are associated with spatial structure of concentration field: organized into puffs or clumps of odor
- In faster flows, agent has less time to "see" this spatial structure

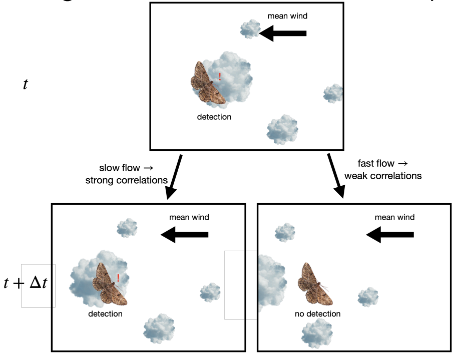

Figure Fast flows decorrelate odor encounters

Correlations, cont'd

- Main question: what is the effect of correlations on the search performance?
- Two possible approaches: (a) ignore correlations, or (b) keep track of previous observation and use $\Pr(o_t | o_{t-1}, \mathbf{s})$
- Strategy:
    1. Run DNS with a source of passive scalars
    2. Tune correlations by rescaling time in flow $t \to \alpha t$
    3. Find quasi-optimal policy with and without correlations
    4. Compare Monte Carlo search performance

Results

The DNS

- 3-D incompressible Navier-Stokes with mean wind U on $1024 \times 512 \times 512$ grid in turbulent regime $\mathrm{Re}_\lambda \simeq 150$
- Periodic BCs, stochastic large-scale forcing
- Lagrangian particles emitted simultaneously from point sources at 5 locations, data dumped every τ_η ($\sim 4000\tau_\eta$ total)
- Have data for 5 different mean flow speeds $U = 0, 1.5, 3, 6, 9$ with $u_{\mathrm{rms}} \approx 1.2$. To our knowledge, only data set of its kind

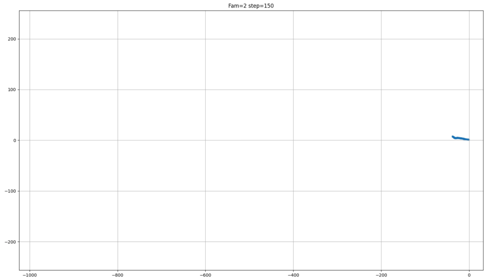

Coarse-graining

- To move to POMDP setting, data are coarse-grained on a quasi-2D slice to obtain 99×33 grid with spacing $\sim 10\eta$
- Grid aspect varied depending on wind speed, fixing total cells
- Particles counted to obtain concentration field

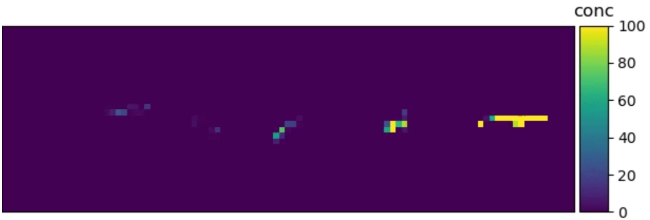

Empirical likelihoods

- Define $c_{\rm thr} \gg \langle c | c > 0 \rangle$
- $\Pr(o|\mathbf{s}) \equiv \Pr(c(\mathbf{s}) \geq c_{\rm thr})$ averaged over time and source locations, symmetrized across wind axis
- Use SARSOP to solve for policy using empirical likelihood

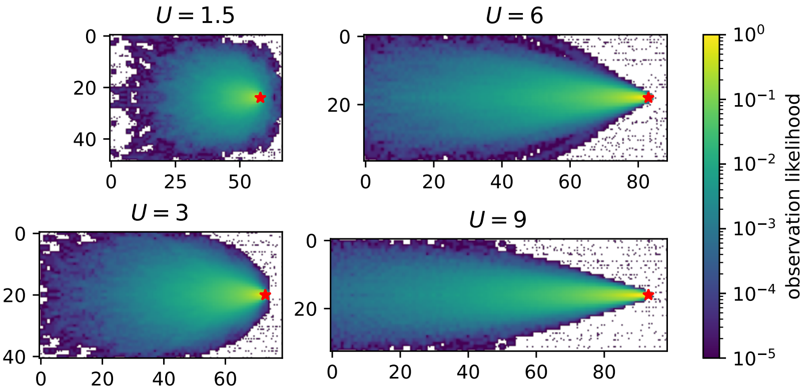

Figure Empirical likelihoods of observation for $c_{\rm thr} = 100$ when $U \neq 0$

Searching in the DNS: near-optimal vs. heuristics

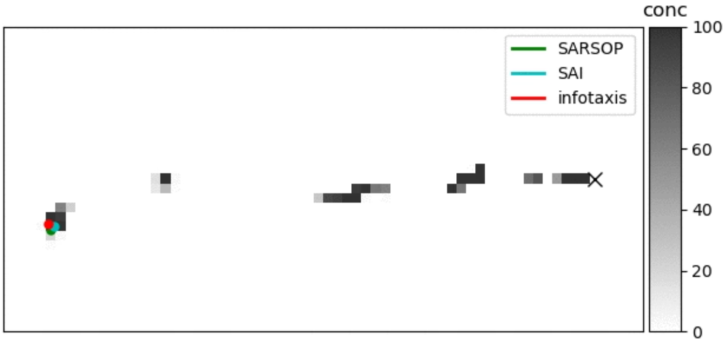

Note casting (crosswind zig-zagging) behavior in all policies! Very similar to real moths

Arrival time statistics for $U = 9$

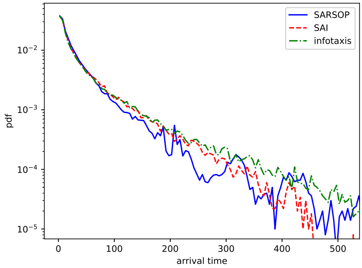

Figure Arrival time pdfs for searching in the source

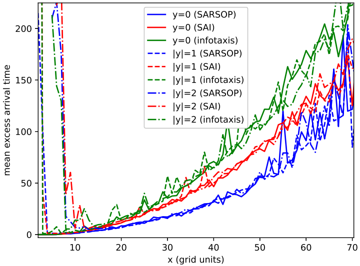

Figure Mean arrival time (minus distance from source) conditioned on starting position

policy	$\mathbb{E}[T \mid T < T_{\max}]$	$\Pr(T \geq 50)$	$\Pr(T \geq 100)$	$\Pr(T \geq T_{\max})$
SARSOP	39.4 ± 0.2	0.223 ± 0.001	0.0951 ± 0.0009	$< 10^{-5}$
SAI	43.0 ± 0.2	0.263 ± 0.001	0.124 ± 0.001	0.0014 ± 0.0001
infotaxis	48.6 ± 0.2	0.277 ± 0.001	0.145 ± 0.001	0.0013 ± 0.0001

Table 1: Arrival time statistics when using the empirical likelihood and searching within the DNS.

Correlation strengths

- Natural measure of strength of correlations is $\delta \equiv p_{11} - p_{10}$. Closely related to correlation time
- $-1 \leq \delta \leq 1$. Sign determines if positively or negatively correlated
- Generally more strongly correlated for small threshold, small U, small time rescaling α

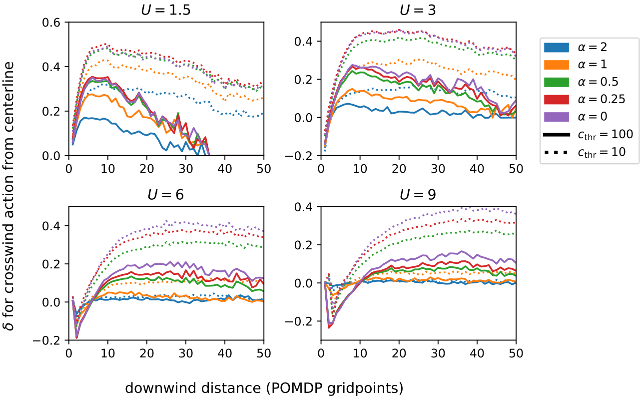

Arrival time performance

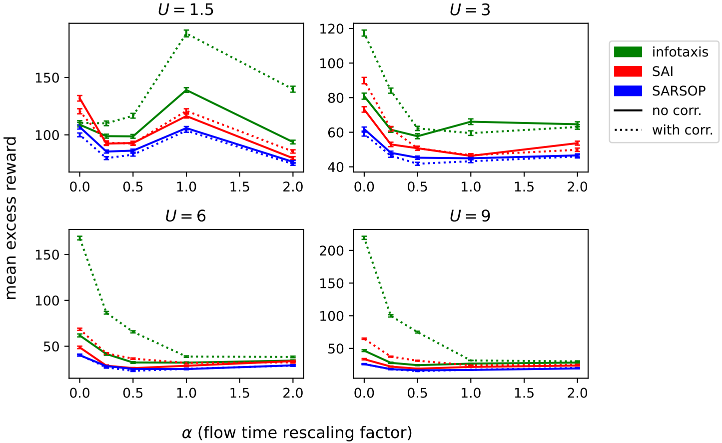

Figure Regularized arrival time performance for $U \neq 0$

Infotaxis is **degraded** by including correlations! If wind is strong, gathering information $\neq$ arriving to the source

Most likely trajectories: quasi-optimal and infotaxis

Most likely trajectory is to detect nothing. Good baseline for understanding policy

$U = 9$

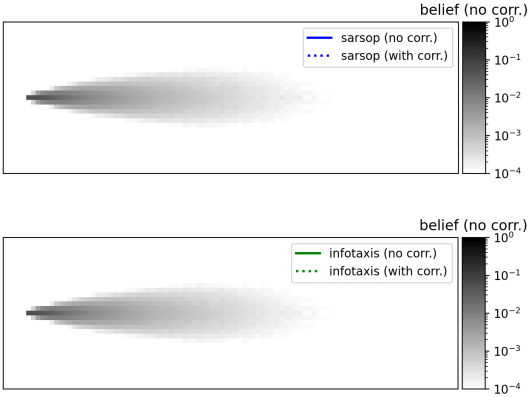

$U = 1.5$

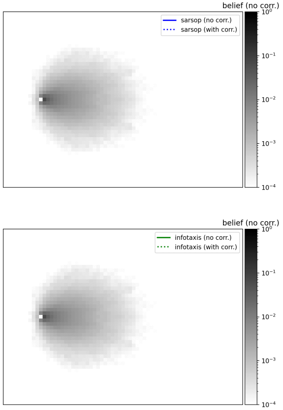

Sketch of theory

- Basic idea: time for posterior to converge in probability depends on δ
- For stationary agent, can show that if agent is unaware of correlations

$$T_{\mathrm{unaware}} - T_0 = \frac{2\delta^* p_0^* p_1^*}{1-\delta^*} \log^2 \frac{p_0}{p_1} = \frac{2\delta^*}{1-\delta^*} T_0,$$

  where * means evaluation at ground truth and no * means evaluation at test point. T_0 is uncorrelated case
- Thus positive (negative) correlations slow down (speed up) the time to estimate the source
- Can also show that taking correlations into account results in a **slight** improvement to convergence ($O(\delta^2)$)
- N.B. full analysis more involved (e.g. can take into account motion of agent, also need to consider time for asymptotic posterior to be informative)

Conclusions

- Generated a high-quality data set for tracers emitted from a point source in a turbulent flow
- Found quasi-optimal policies to search in the flow, with and without correlations. When mean wind is sufficiently strong, optimal motion is to cast à la real moths
- Strong correlations degrade search performance by slowing convergence of posterior
- Infotaxis **fails** when correlations are included in the policy (and mean wind is sufficiently strong)
- Showed results for $U \neq 0$ today. Analysis of isotropic case not yet complete

References I

Celani, A., Villermaux, E., and Vergassola, M. (2014).
Odor landscapes in turbulent environments.
Physical Review X, 4(4):041015.

Loisy, A. and Eloy, C. (2022).
Searching for a source without gradients: how good is infotaxis and how to beat it.
Proceedings of the Royal Society A, 478(2262):20220118.

Vergassola, M., Villermaux, E., and Shraiman, B. I. (2007).
'Infotaxis' as a strategy for searching without gradients.
Nature, 445(7126):406–409.

Villermaux, E. and Innocenti, C. (1999).
On the geometry of turbulent mixing.
Journal of Fluid Mechanics, 393:123–147.