Belief Change Maximisation for Hydrothermal Vent Hunting using Occupancy Grids

Zeyn Saigol[*], Richard Dearden[*],
Jeremy Wyatt[*] and Bramley Murton[†]

[*]School of Computer Science
University of Birmingham

[†]National Oceanography Centre
Southampton

TAROS 2010, Plymouth

Outline

- Motivation – vent prospecting
- Problem details
- Original algorithms
  - Single-step lookahead: Entropy and $\Sigma\Delta H$
  - Non-myopic planning: $\Sigma\Delta H$-MDP
  - Fix for re-rewarding: OP correction
- Summary

Motivation – Hydrothermal Vents

- Sea floor, 3000m, 350°C
- Emit a plume containing ‘tracers’, dissolved chemicals and minerals
- Turbulent current means no gradient
- Often found in clusters, so plumes combine

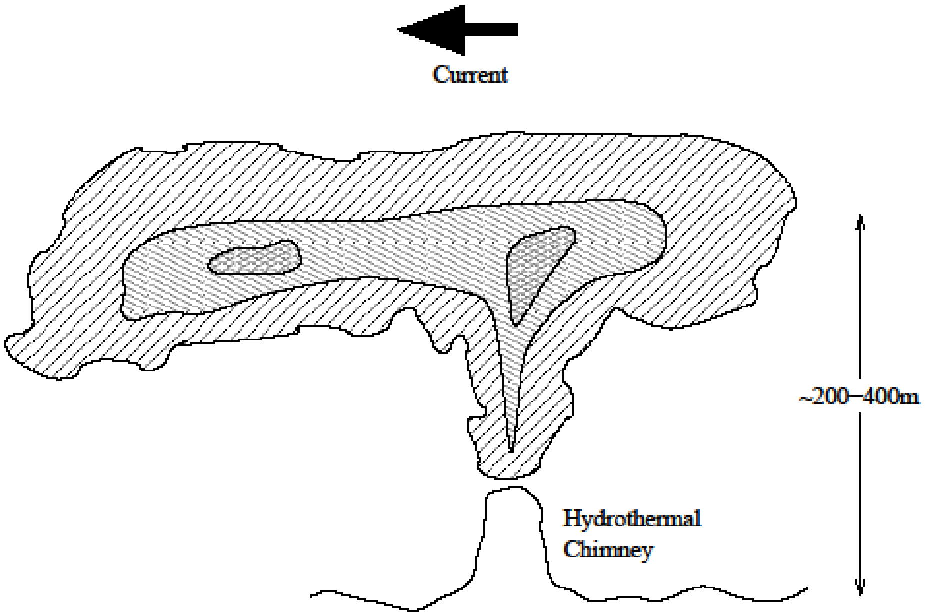

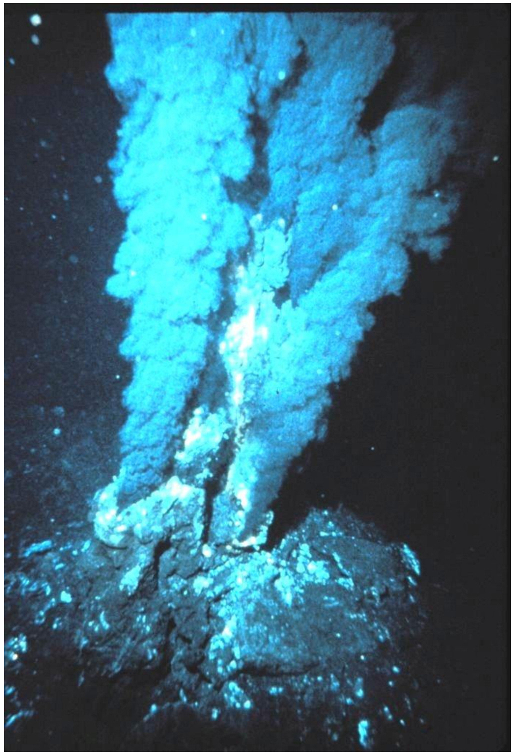

The Challenge

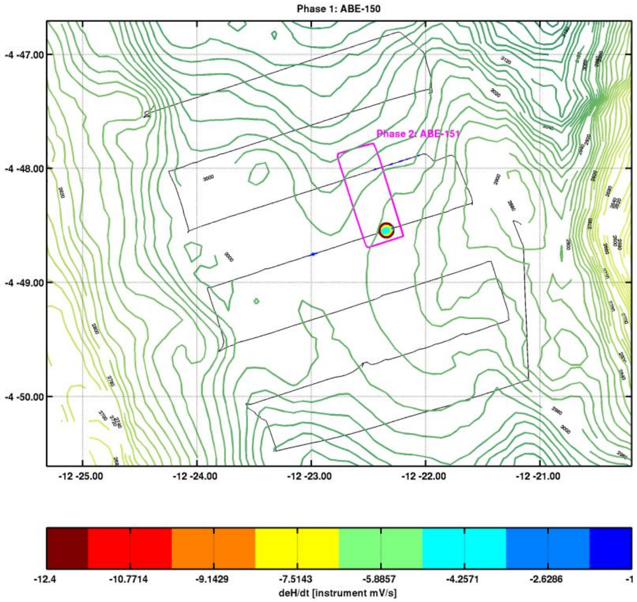

- Ship-based search followed by AUV deployment
- Use chemical tracers – vision impossible, sonar difficult
- AUVs – exhaustive search
- Use AI: goal of finding as many vents as possible during mission
- Partially observable, multiple sources, indirect observations
- Options:
  - Reactive, moth-like
  - Information theoretic – build probabilistic map, then plan

Outline

- Motivation – vent prospecting
- Problem details
- Original algorithms
  - Single-step lookahead: Entropy and ΣΔH
  - Non-myopic planning: ΣΔH-MDP
  - Fix for re-rewarding: OP correction
- Summary

Problem Model

- Mapping: adopt occupancy grid (OG) algorithm of Michael Jakuba
- Uses plume detections and current to infer map. Observations $z \in$ {locate vent, detect plume, nothing}
- Cells occupied ($m_c \Leftrightarrow$vent) or empty; OG consists of $P(m_c)$ values
- Belief state b = (OG, $\mathbf{x}_{AUV}$)
- Actions: $a \in \{N,E,S,W\}$
- OG : b'=`srog`(b,a,z)
- Observation model P(z|b,a)
- Partially-observable Markov decision process (POMDP) – but intractable, 20x20 grid => 10^{244} states

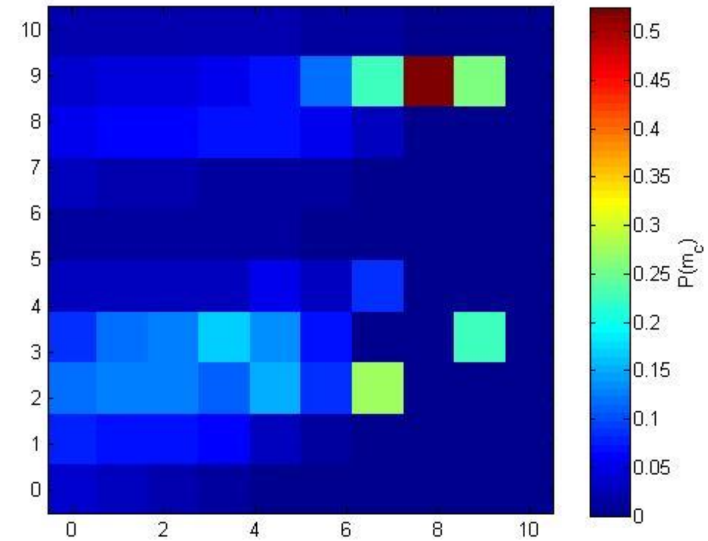

Outline

- Motivation – vent prospecting
- Problem details
- Original algorithms
  - Single-step lookahead: Entropy and $\Sigma\Delta H$
  - Non-myopic planning: $\Sigma\Delta H$-MDP
  - Fix for re-rewarding: OP correction
- Summary

Infotaxis Algorithm

- Vergassola *et al.* developed infotaxis for finding a single chemical source, using a continuous distribution map
- Chooses action that reduces uncertainty in map the most
- Uncertainty defined by entropy; entropy of OG from sum of entropy of cells H_c

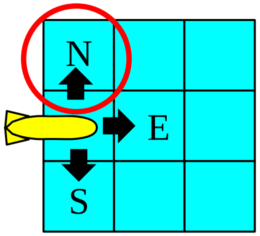

Infotaxis Algorithm

- Vergassola *et al.* developed infotaxis for finding a single chemical source, using a continuous distribution map
- Chooses action that reduces uncertainty in map the most
- Uncertainty defined by entropy; entropy of OG from sum of entropy of cells H_c

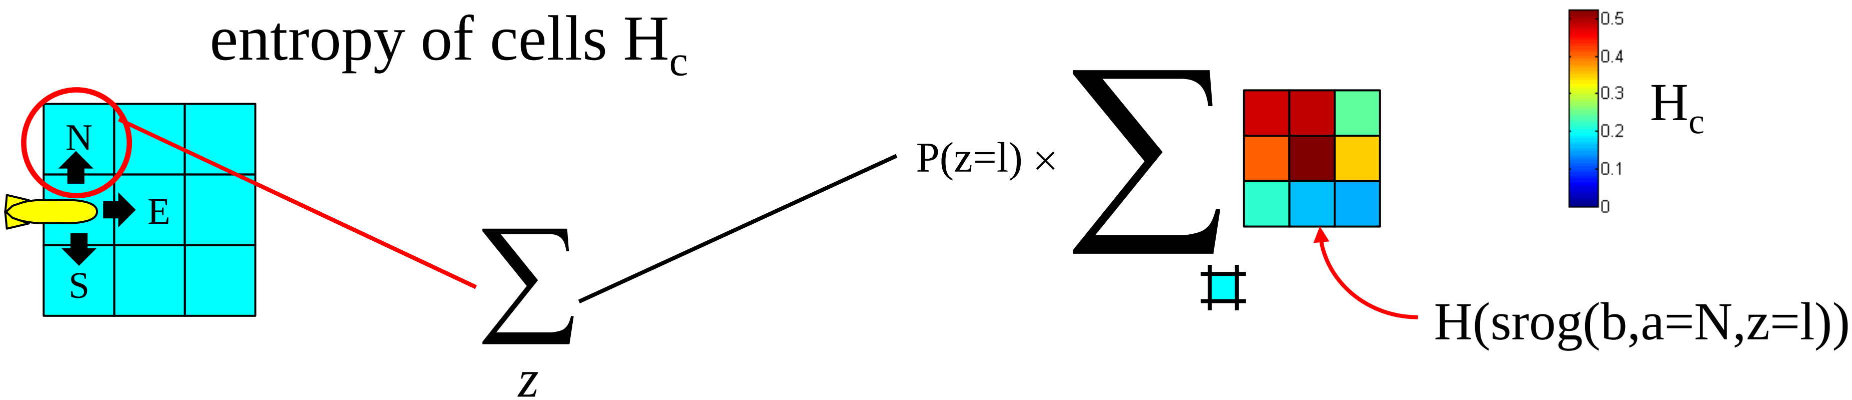

- Value of N action = expected new entropy

- Entropy for observation *locate vent*

Infotaxis Algorithm

- Vergassola *et al.* developed infotaxis for finding a single chemical source, using a continuous distribution map
- Chooses action that reduces uncertainty in map the most
- Uncertainty defined by entropy; entropy of OG from sum of entropy of cells H_c

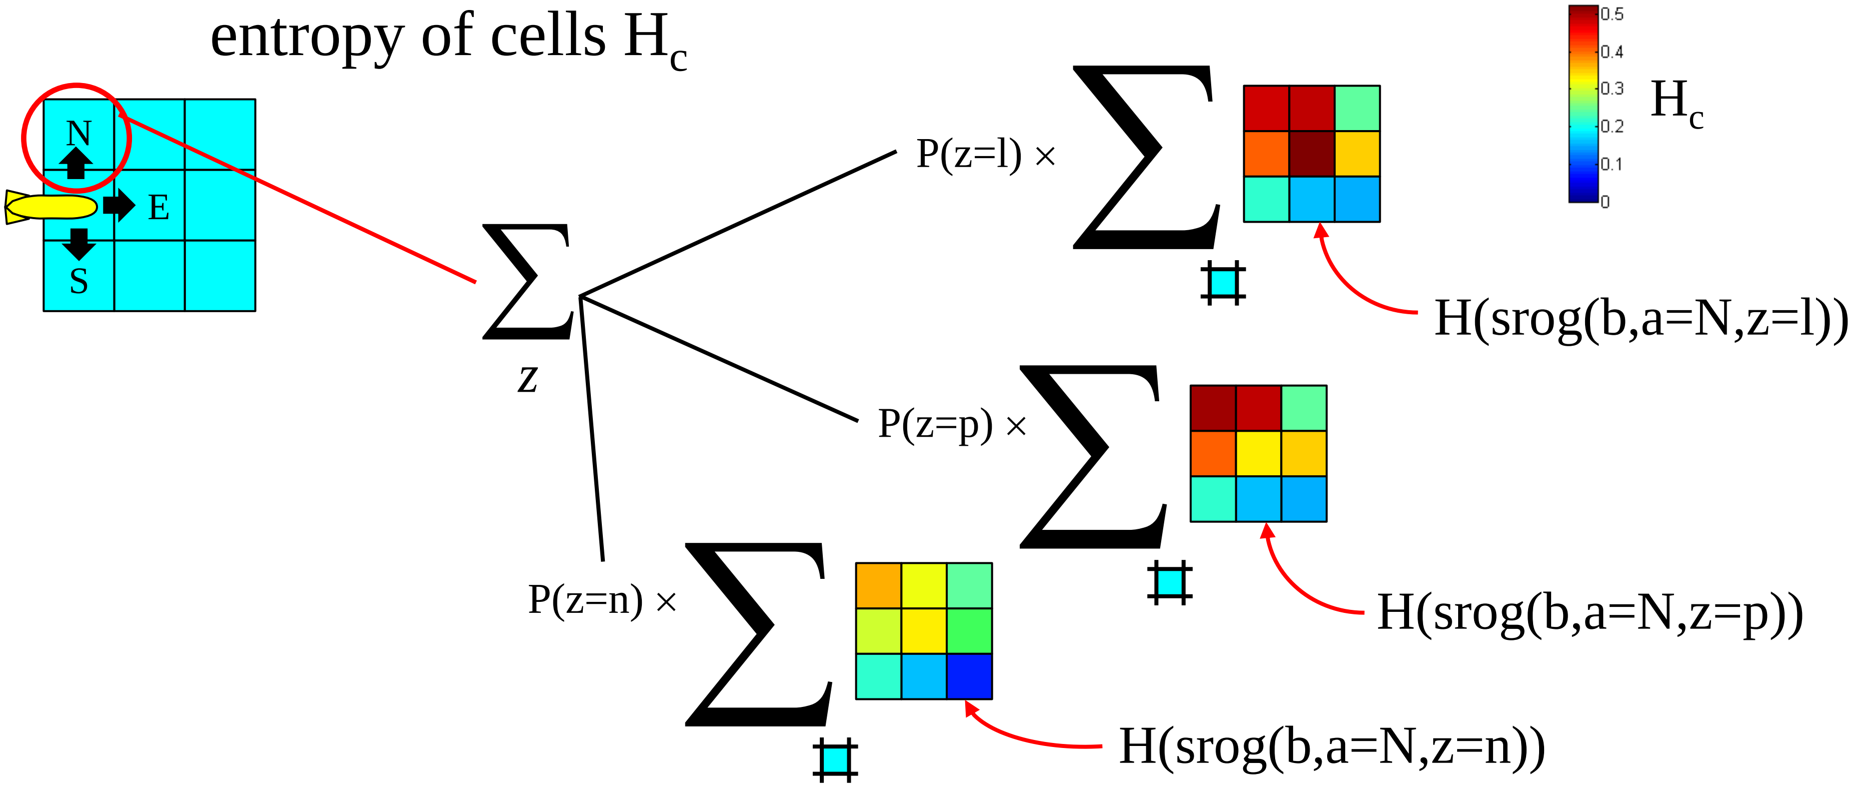

Infotaxis Algorithm

- Vergassola *et al.* developed infotaxis for finding a single chemical source, using a continuous distribution map
- Chooses action that reduces uncertainty in map the most
- Uncertainty defined by entropy; entropy of OG from sum of entropy of cells H_c

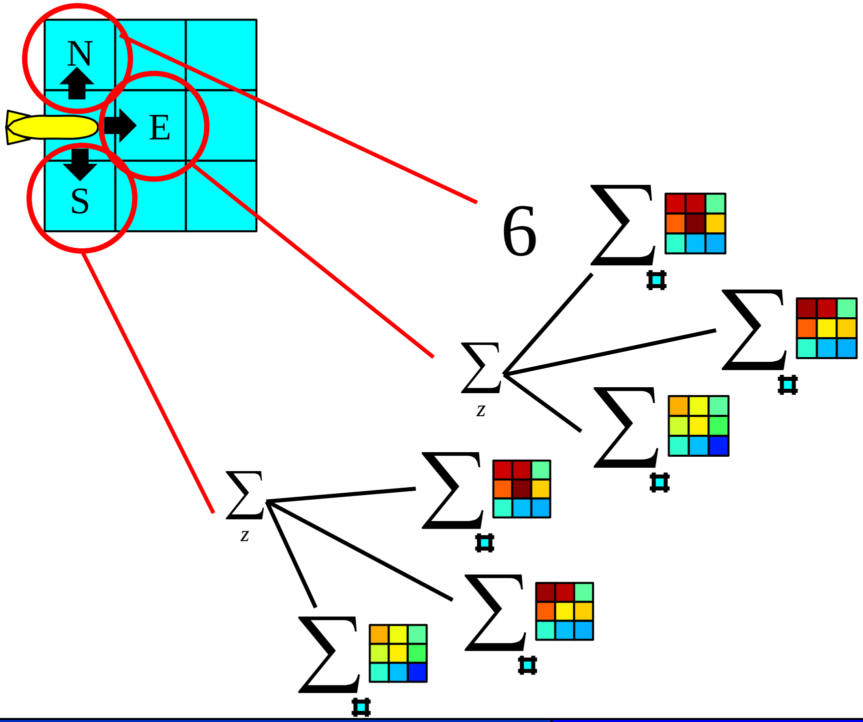

Infotaxis Algorithm

- Vergassola *et al.* developed infotaxis for finding a single chemical source, using a continuous distribution map
- Chooses action that reduces uncertainty in map the most
- Uncertainty defined by entropy; entropy of OG from sum of entropy of cells H_c

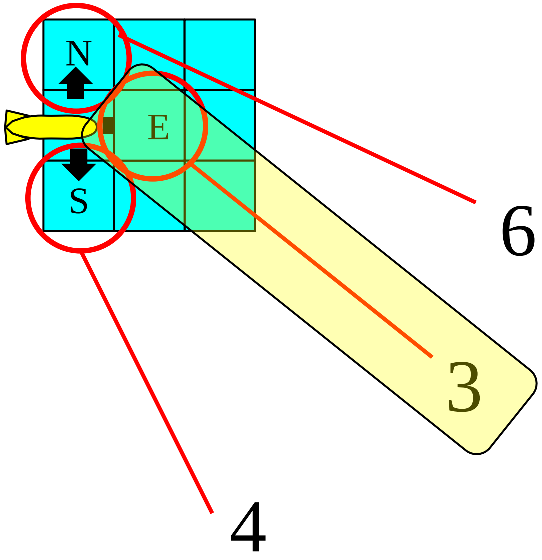

ΣΔH Algorithm

- Jakuba's OG algorithm requires a low prior occupancy probability (0.01), as small number of vents are expected in a given search area
- This means plume and vent detections, which provide useful information, can actually increase entropy

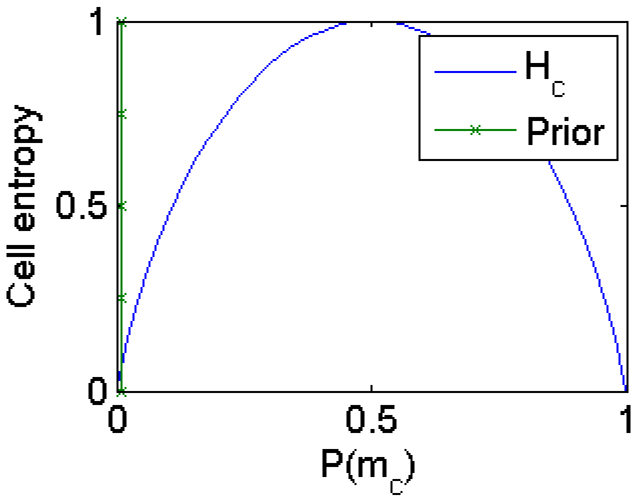

ΣΔH Algorithm

- Jakuba's OG algorithm requires a low prior occupancy probability (0.01), as small number of vents are expected in a given search area
- This means plume and vent detections, which provide useful information, can actually increase entropy
- Heuristic alternative: ΣΔH. Use the *change* in entropy, regardless of whether increase or decrease

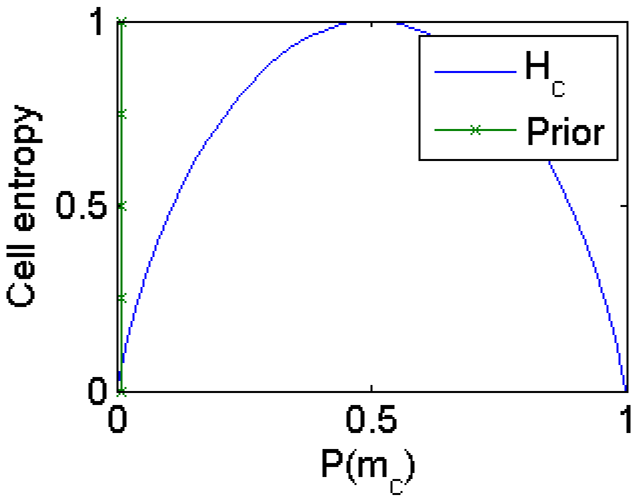

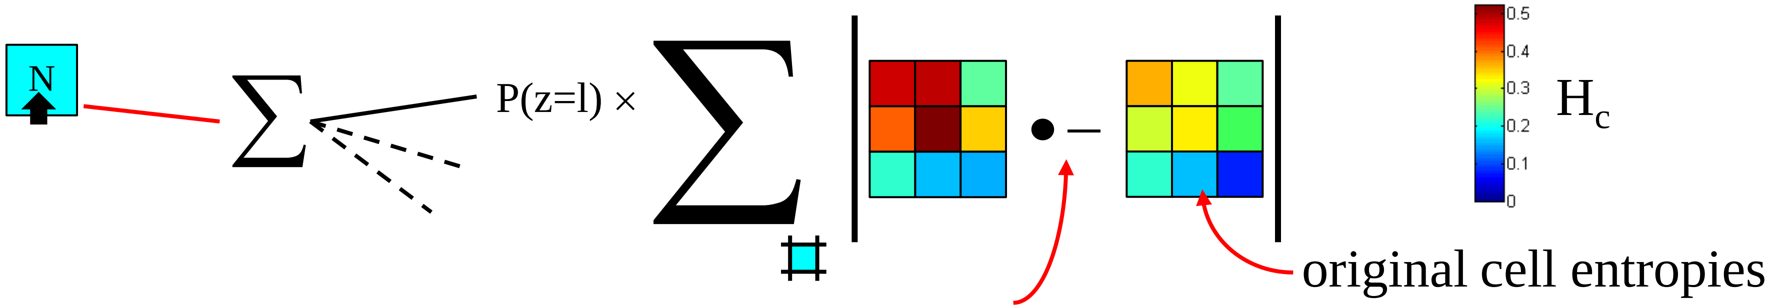

Outline

- Motivation – vent prospecting
- Problem details
- Original algorithms
  - Single-step lookahead: Entropy and $\Sigma\Delta H$
  - Non-myopic planning: $\Sigma\Delta H$-MDP
  - Fix for re-rewarding: OP correction
- Summary

Non-myopic Planning: ΣΔH-MDP

- Issue: only plan one step into future
- Intuition: instead of evaluating possible action/observation pairs N steps into future, evaluate effects of observations N steps away – avoids exponential blowup

Non-myopic Planning: ΣΔH-MDP

- Issue: only plan one step into future
- Intuition: instead of evaluating possible action/observation pairs N steps into future, evaluate effects of observations N steps away – avoids exponential blowup
- Mechanics:
  - Calculate $E_z[\Sigma\Delta H]$ for making an observation from a cell, for *every* cell in the OG
    (as if AUV could teleport to any cell)
  - Assume that the OG no longer changes, and define a reward of $E_z[\Sigma\Delta H]$ for visiting a cell
  - Then solve a deterministic Markov decision process (MDP) to get the optimal policy given these assumptions

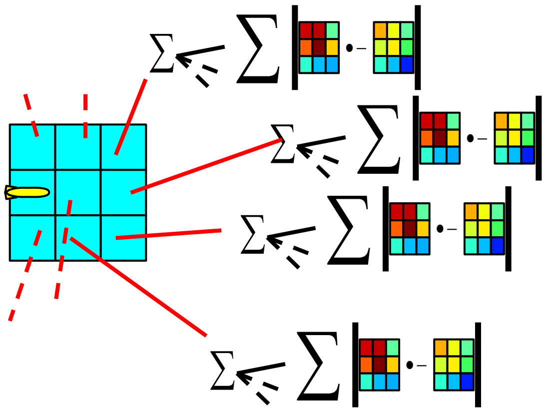

ΣΔH-MDP Movie

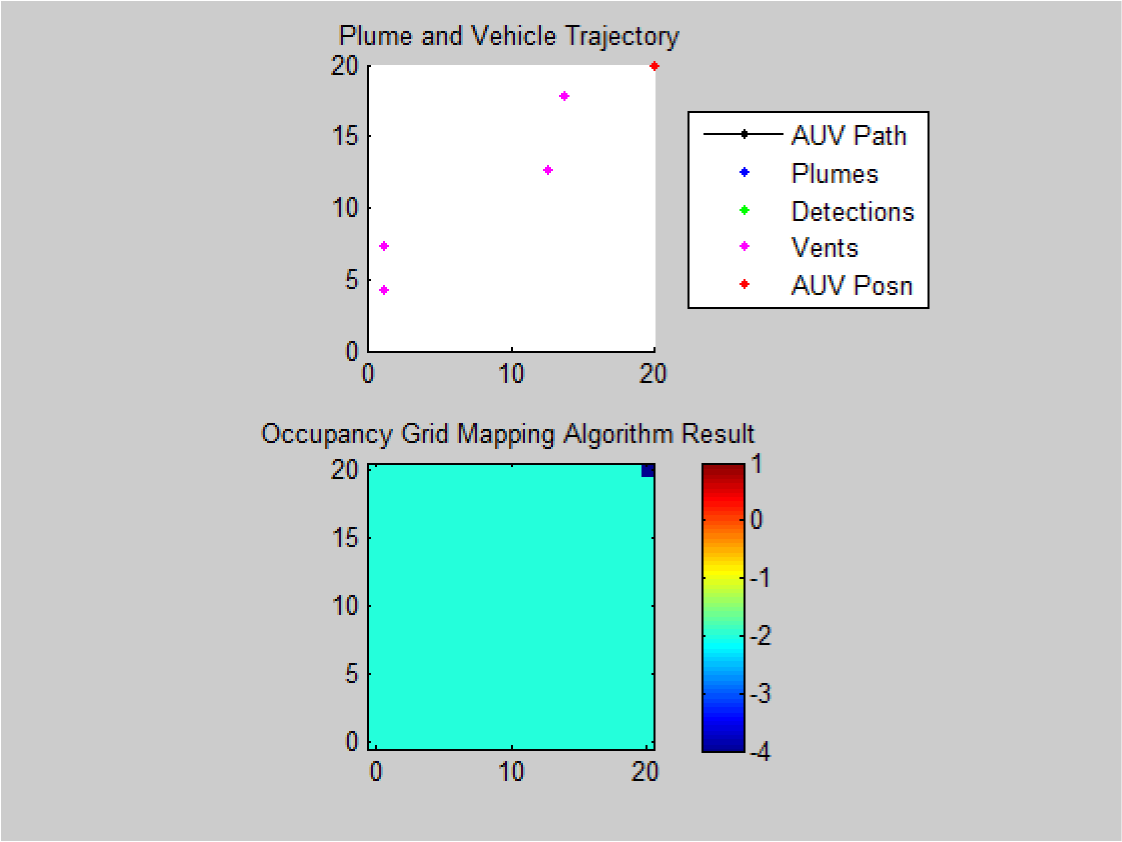

Results

- Setup: percent found, 133 timesteps, mean of 600 trials
- ΣΔH significantly better than mowing-the-lawn (MTL)

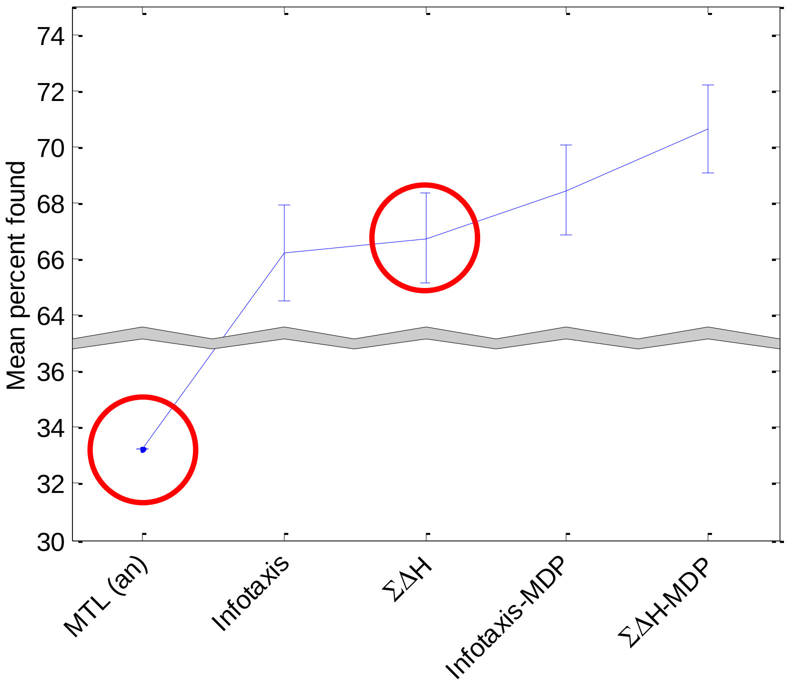

Results shown with 95% confidence intervals

Results

- Setup: percent found, 133 timesteps, mean of 600 trials
- ΣΔH significantly better than mowing-the-lawn (MTL)
- ΣΔH-MDP improves on ΣΔH

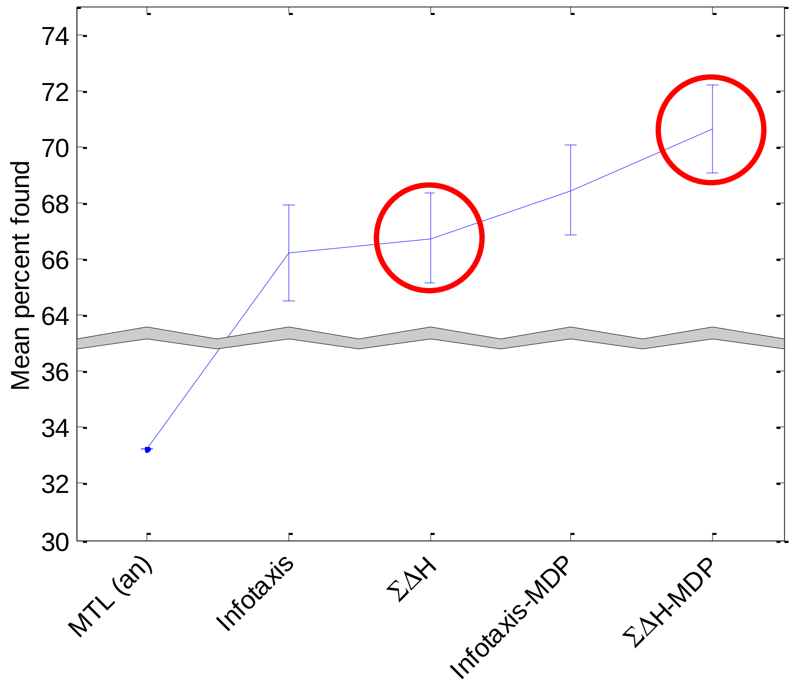

Results shown with 95% confidence intervals

Results

- Setup: percent found, 133 timesteps, mean of 600 trials
- ΣΔH significantly better than mowing-the-lawn (MTL)
- ΣΔH-MDP improves on ΣΔH
- ΣΔH improves on infotaxis

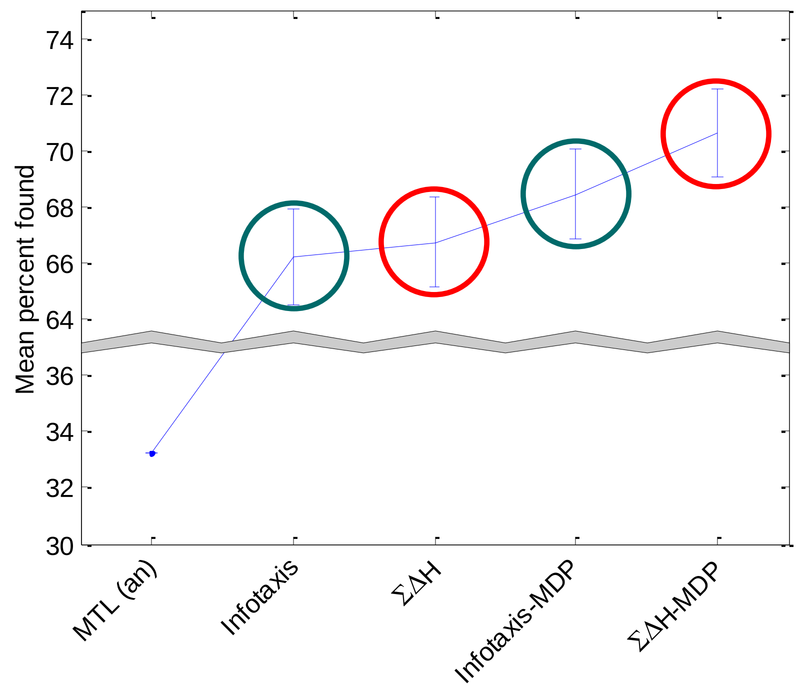

Results shown with 95% confidence intervals

Outline

- Motivation – vent prospecting
- Problem details
- Original algorithms
  - Single-step lookahead: Entropy and $\Sigma\Delta H$
  - Non-myopic planning: $\Sigma\Delta H$-MDP
  - Fix for re-rewarding: OP correction
- Summary

OP Correction

- Slight issue with ΣΔH-MDP is that the MDP assumes re-visiting a cell earns the same reward
- In fact, repeated observations from same cell are worth less
- ΣΔH-OP: replace the MDP with an *Orienteering Problem* solver
- Flag-gathering task – zero reward for re-visiting a cell
- OP is a variant of the TSP with rewards for cities and a limited path length
- Use a Monte-Carlo method: generate random non-crossing paths and select the best

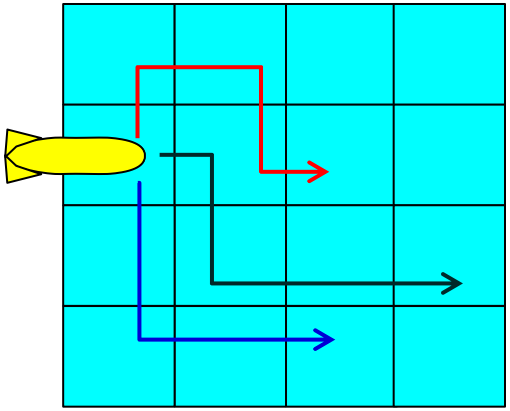

Results - OP Correction

- Results compared to IL4 – online POMDP – our previous state-of-the-art solution for this domain (Saigol *et al.* 2009)
- Also applied to OP correction to IL with less conclusive results (see paper)

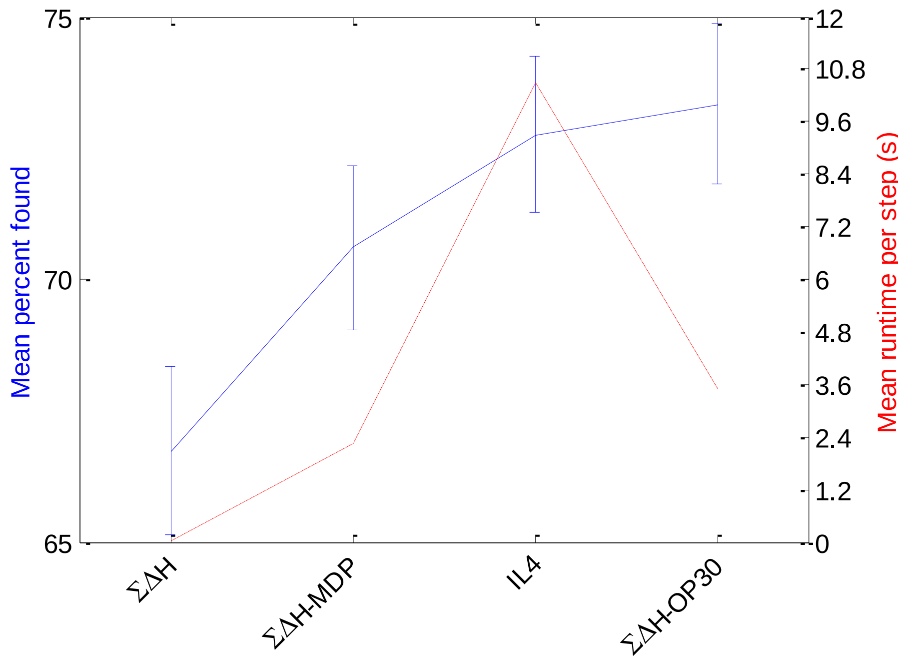

Summary

- We have formalised an interesting real-world problem that poses a significant challenge for AI
- We have created a novel ΣΔH-MDP algorithm to guide exploration in occupancy grids
- This adapts existing entropy-based techniques to deal with:
  - Low prior occupancy probabilities
  - Uncertain, long-range sensors
  - Planning further into the future
- When an OP correction is applied, ΣΔH-OP significantly outperforms traditional methods such as MTL, and performs at least as well as online POMDP methods but requires less computation time

References

Jakuba, M. (2007). *Stochastic Mapping for Chemical Plume Source Localization with Application to Autonomous Hydrothermal Vent Discovery*. PhD thesis, MIT and WHOI Joint Program.

Saigol, Z., Dearden, R., Wyatt, J., and Murton, B. (2009). Information-lookahead planning for AUV mapping. *Proceedings of the Twenty-first International Joint Conference on Artificial Intelligence (IJCAI-09)*.

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

Questions

- Any questions?