

Improving the robustness of Naïve Physics airflow mapping, using Bayesian reasoning on a multiple hypothesis tree. *

Gideon Kowadlo
Intelligent Robotics Research Centre
Monash University
Clayton, Victoria, Australia
gkowadlo@ieee.org

R.Andrew Russell
Intelligent Robotics Research Centre
Monash University
Clayton, Victoria, Australia
andy.russell@eng.monash.edu

Abstract—Previous work on odour localisation in enclosed environments, relying on an airflow model, has faced significant limitations due to the fact that large differences between airflow topologies are predicted for only small variations in a physical map. This is due to uncertainties in the map and approximations in the modelling process. Furthermore, there are uncertainties regarding the flow direction through inlet/outlet ducts. We have presented a method for dealing with these uncertainties, by generating multiple airflow hypotheses. As the robot performs odour localisation, airflow in the environment is measured and used to adjust the confidences of the hypotheses using Bayesian inference. The best hypothesis is then selected, which allows the completion of the localisation task. We have shown experimentally that this method is capable of improving the robustness of our method for odour localisation in the presence of uncertainties, where previously it was incapable. The results further demonstrate the usefulness of Naïve Physics for practical robotics applications.

Index Terms—Mapping, Naïve Physics, Odour Localisation, Bayesian, Multiple Hypothesis.

I. INTRODUCTION

Odour localisation is performed by various organisms, and is now a burgeoning area of robotics research. Work in this field, and biorobotics in general, contributes to knowledge of biological organisms, and inspires innovation and new applications. Robotic odour localisation experiments have explored and mimicked the chemotactic behaviour of such organisms as the silk worm moth (locate mates) [1], lobsters (find food) [2], dung beetles (find feces for brooding, living in or eating) and *e-coli* bacteria (find nutrients) [3]. Potential applications for odour localising robots include finding humans in search and rescue operations, locating gas leaks in industry, finding fires in their initial stages as well as many other possibilities.

In general, odour particles are carried downwind, forming a plume that spreads, meanders and becomes patchy. Gradients of concentration and airflow are exploited using reactive behaviour by various organisms, and robots [3]–[5]. In some enclosed cluttered environments (low ceiling and thinly populated by objects that affect airflow) the plume is poorly defined, and odour dispersal is dominated by the formation of sectors. Odour can be confined to sectors (local or downwind

from the release location). The reactive behaviour of previous solutions is not suited to, and therefore unreliable in these types of environments [6], which may be encountered in a cave, air duct, sewer or crawl-way beneath a house.

In previous work, we have shown that the problem can be tackled effectively using a ‘sense-map-plan-act’ style control strategy [7], [8], comprised of firstly mapping the airflow in the environment using a Naïve Physics algorithm (NaReM) [9], and then using this map to reason about odour dispersal, plan an information gathering stage, navigate the environment taking readings, and finally make a prediction. Naïve Physics is the chosen method for airflow modelling as it avoids many of the difficulties of using computational fluid dynamics, and provides data structures at a high level of abstraction, readily used by the reasoning algorithm. In this project, the odour source location prediction was restricted to a physical domain. In order to make the localisation specific, vision was combined with olfaction to perform a bi-modal complementary sensing search [10].

For this approach to work effectively, the topology of sectors (from the airflow map), must be correct. This is very difficult to achieve, as small variations in the environment and in the implementation of the airflow modelling, can result in extremely different, but stable, airflow patterns. Furthermore, there are significant uncertainties in the information available to the robot prior to computing the airflow model, arising from the physical map provided to the robot, boundary conditions, and the application of the airflow modelling rules.

Uncertainty leads to problems of data association, which have been tackled extensively in the areas of tracking, localisation and mapping, with the use of Kalman filters, particle filters, and other methods based on Bayesian reasoning. Further, multiple hypothesis trees have been used since the late 60’s (see [11] for a good review), and combined with Bayesian reasoning by Reid [11]. Similar methods have been used for tracking [12], as well as for a variety of other applications from genetics (phylogeny) [13], to image processing (edge grouping and contour segmentation) [14], and mapping (modelling a dynamic environment) [15].

In this paper, we demonstrate an effective method for improving the robustness of the approach, in the presence of imprecise information and uncertainties. Where uncertainties

*This work is supported by the ARC funded Centre for Perceptive and Intelligent Machines in Complex Environments.

exist, new hypotheses are spawned and contained in a tree that retains important relational information. Each hypothesis has a confidence level, which is continually adjusted. The robot navigates the environment using a modified version of the previous, single hypothesis approach. As it moves, it records the direction of airflow, and identifies locations for which comparisons between predicted and measured airflow can be used as evidence for hypotheses. Bayesian inference is then used to adjust the confidence levels of the set of hypotheses. Finally, the best hypothesis is chosen, enabling the prediction of the odour source location. The method is robust to bad evidence, erroneous measurements, and imprecise sensors.

Multiple hypothesis trees are conventionally used with an iterative Bayes filter, involving a continually growing tree. Bayesian reasoning is used to modify the conditional probabilities of nodes, which become the prior probabilities for child nodes, and this is used to prune branches. Our methodology differs significantly, in that the tree is fully determined prior to taking measurements, and then the probabilities of nodes are continually adjusted.

This paper includes a description of the uncertainties in the airflow mapping (Section II), the methods for generating multiple airflow map hypotheses (III) and resolving these hypotheses during and for the odour localisation task (IV), as well as the experimental setup (V) and results (VI).

II. UNCERTAINTIES

A. Sources of uncertainty

During the airflow mapping, there are three major sources of uncertainty:

a) Uncertainty in the physical map: In most robotics applications, the map obtained by or provided to the robot will be imprecise. In this implementation, an *a priori* physical map is provided. For many applications, such as for gas leak detection in a known industrial area, such a map is likely to be available.

b) Unknown boundary conditions: In many cases, even with a highly certain physical map, the direction of airflow through a duct may be unknown. Under the conditions described for the target environment, this is the most significant effect of unknown boundary conditions. Other parameters such as the temperature of walls will have a negligible effect.

c) Naïve Rules: Due to the nature of the problem, in which there is a formation of distinct stable airflow patterns, the algorithm has not been designed to combine rules to varying degrees. Although in many cases, a rule will be applied if a flow parameter is within a certain range (i.e. flow angle when encountering a solid surface), the edge of the range is a hard decision boundary. However, at the edge of decision boundaries, more than one course of action could be reasonably predicted; often leading to large macroscopic flow pattern discrepancies.

A second issue arises from the use of a sequential algorithm, to model a physically concurrent process. Airflows spread out laterally as they move forward. However, the algorithm traces the flows in layers; the inlet flow is propagated




Fig. 1. Flow patterns.

as a discrete flow, until it is complete, and then further layers are initiated around this flow to simulate the spreading of the flow laterally. The possibility of incorporating progressive lateral spreading of the flow has been considered, however, the degree of spreading is unknown and gives rise to many further uncertainties that are difficult to manage. In addition, it would not account for non-uniform certainty in the physical map, and would only predict one possible flow pattern (which may still be inaccurate).

The current flow propagation method is an effective way of tracing out sectors that define a stable formation. Combined with the multiple hypothesis scheme described in this paper, it is more robust, being able to accommodate uncertainty in the physical map, and the large imprecision of the naïve model; as well as providing alternative hypotheses.

B. Decision points

The uncertainties manifest as two ‘decision points’ (analogous to a data association uncertainty in a problem such as tracking) in the airflow modelling:

1. Air duct flow direction: In the case of unknown boundary conditions.
2. Flows past obstacles: The scenario illustrated in Fig. 1, shows the flow patterns that arise with differing vertical positions of two types of obstacles (simulated with Flo++ and verified through practical experimentation [6]). Significantly different, stable, macroscopic flow patterns occur with even small changes in the vertical position of the obstruction. This is due to the flow spreading out as it propagates, being partially redirected by obstacles, and the degree to which these occur - which is difficult to predict and depends on many factors. Therefore, due to uncertainties in the physical map (position of obstacles), and the NaReM (position of flows and implementation of rules), the question of redirection of a flow as it moves past an obstacle constitutes a decision point.
3. Others: Other decision points include: ‘are flows oncoming?’, ‘are flows perpendicular (to each other, solid/wall)?’, and the case of the obstacle causing complete (as opposed to partial) redirection of the oncoming flow.

III. GENERATION OF MULTIPLE HYPOTHESIS TREE

In the case of the two main decision points (explained above), more than one outcome is possible. For each possibility, new hypotheses are spawned, and contained in a tree.

A. Air duct flow direction

It is assumed that the airflow direction through the ducts is dependent on conditions that are unknown *a priori*. In addition, the robot begins from within the environment. Alternatively, if the robot moved from outside to within the enclosed area, it would be possible to measure the airflow as it passes the threshold (there are potentially other cues as well, such as local airflow direction). However, as the direction is unknown, multiple hypotheses are formed for each combination of inlet/outlet, leading to several sibling root nodes. A probability is assigned to each root node. In this case equal probability is assumed, however in some circumstances, partial knowledge may be available or deduced using cues.

B. Flows past obstacles

As seen in Fig. 1, and described above, the flow may be partially redirected parallel to the obstacle or just continue. Therefore, the area surrounding the leading edge of the flow is examined. If there is an obstacle in the circular vicinity, for which it is ambiguous whether flow will be diverted or not, and if a new hypothesis has not already been initiated for this flow/object pair, then the current hypothesis is frozen, and two new hypotheses (one with flow redirected, one with it flowing past) are initiated and added as children to the current hypothesis. The latter is a copy of the current hypothesis (as the flow was not in the immediate vicinity, and would have continued unimpeded). In order to qualify as an ambiguous obstacle, it must be in the ‘non-immediate’ vicinity, 3.5-10cm, the surface of the obstacle is mostly perpendicular to the flow, and the flow is moving towards this obstacle, but an intersection is not imminent. In addition, the flow must be of sufficient strength, in order to be redirected. The same criteria are used for the identification of salient information, and explained in Section IV-A.1.

Probabilities are assigned to the new hypotheses. In this case, they are again assumed to have equal probability. However, physical map certainty (which would be available if a statistical mapping method is employed) could be used to assign more appropriate values. When the hypothesis is added to the tree, the final value of the ‘node’ probability is the ‘hypothesis’ probability multiplied by its parent’s probability. Therefore, all siblings have probabilities summing to their parent’s probability.

C. Hypothesis tree

Each root node is a different possibility for air duct flow direction. Every subsequent node that has children is a decision point, with its children being the possible outcomes. The leaf nodes are a suite of complete, viable, hypotheses, with probabilities summing to unity. Probabilities are interpreted as confidence in a hypothesis. Let $\Omega \triangleq \{\Omega_i, i = 1, 2, \dots, I\}$ denote

the set of all hypotheses, and $\Omega^L \triangleq \{\Omega_m^L, m = 1, 2, \dots, M\}$ the set of leaf nodes. An example of the hypothesis tree is shown in Fig. 6.

IV. RESOLUTION OF MULTIPLE AIRFLOW HYPOTHESES

The best hypothesis is selected for, and in the context of, the odour localisation task. An overview of the system is shown in Fig. 2. The modules regarding multiple hypotheses are explained in the following sections. Details for the other modules, adapted versions of the single hypothesis reasoning algorithm, can be found in [8].

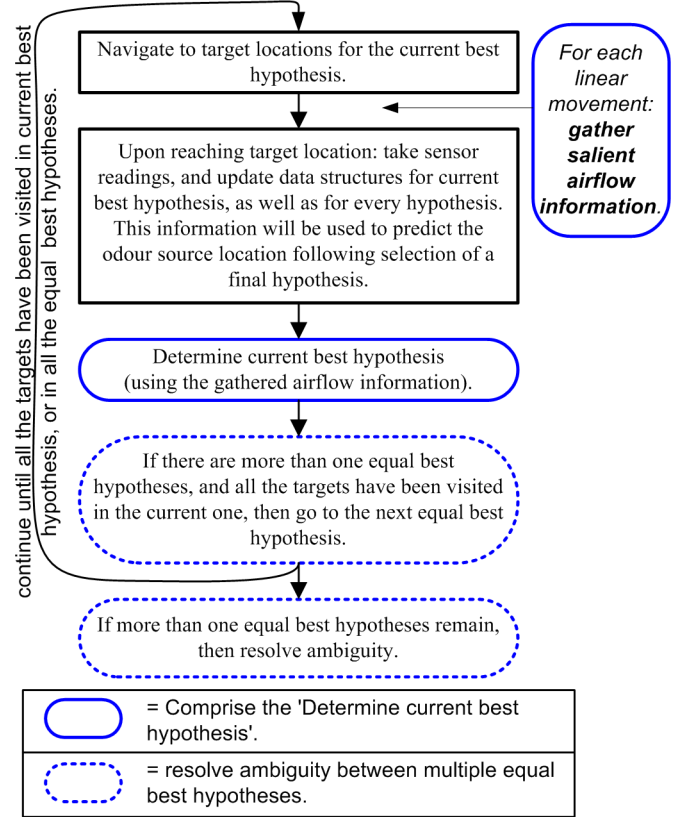


Fig. 2. Resolving multiple hypotheses.

A. Determine current best hypothesis

This section describes the use of evidence to modify the confidence values of the hypotheses. It forms the core of the approach.

Ω^L have associated confidences, which comprise the prior probabilities $p(\Omega_m^L)$. At the end of each robot movement, the path and recorded airflow measurements are analysed for each hypothesis in Ω^L . Viable measurement locations are identified. Each one is used to construct an Evidence for the hypothesis in which it was identified, that may or may not be observed. The Evidence comprises a location, the predicted wind direction, a vector of hypotheses in which it occurs, and a flag to indicate if observed. Many of these points occur over a given area, but comprise one ‘piece’ of information. Therefore, they are the set of preliminary

evidences $Z \triangleq \{Z_n, n = 1, 2, \dots, N\}$, and are clustered into a set of evidences $E \triangleq \{E_k, k = 1, 2, \dots, K\}$. The conditional probabilities $p(E_k|\Omega_m^L)$ are determined, and Bayesian inference is used to calculate the posterior probabilities of $p(\Omega_m^L|E_k)$, with the evidences that have been observed. It is then possible to select the best hypothesis. *Note:* As the robot conducts the odour localisation task, Z grows, and E is repeatedly recomputed. Each of these stages is explained.

1) *Create set of preliminary evidences Z :* In order to afford salient information, the flow must be of sufficient strength to be sensed reliably. The sector boundaries are most appropriate, as they are strong, as well as being a distinctive feature of the hypothesis. In addition, flows that have originated from the inlet, or surround such a flow (and therefore have been initiated as a spreading out of this flow) are considered to be strong. Therefore, for each movement, if the path intersected a ‘strong’ flow, in any of the hypotheses, then the intersection is a salient point that affords a preliminary potential evidence, Z_n , that is added to Z .

2) *Create set of evidences E :* Each time a movement is completed, Z are clustered using a modified k-means algorithm. It determines clusters of compatible preliminary evidences (predicted airflow direction to within 45°), with the assumption of an elliptic distribution (as opposed to circular, as for the conventional k-means). The cluster centres are a representative evidence of their respective cluster (mean position and airflow direction), and comprise E . Similar evidences in E - matching location (within 10cm) and wind direction (45°) - are combined into a single evidence, but associated to the hypotheses from which each of these evidences was identified.

Evidence is considered to be observed if the direction of predicted and measured airflows are within 45° . The subset of observed evidences is denoted by $E^{Obs} \triangleq \{E^{Obs}_j, j = 1, 2, \dots, J\}$. The measured wind direction is calculated by recording a log of wind bearing as the robot travels, every 5cm. The wind values along the travelled line, within a distance of 10cm (hence 20cm linear region) of the evidence location, are averaged with a circular mean function. The predicted airflow is calculated by taking the average of all the flow segments of the model in the circular vicinity of radius 7cm centred about the intersection point.

Fig. 3 shows preliminary evidences Z and post clustered evidences E .

3) *Determine conditional probabilities:* For a given evidence and hypothesis pair E_k and Ω_m^L , determination of the conditional probability $p(E_k|\Omega_m^L)$ is carried out with Algorithm 1. This is done for every pair. Conflicting evidences have the same location (to within 15cm, but airflow direction disparity of greater than 45°). Corresponding evidences are within 15cm of each other.

4) *Update hypothesis confidences:* The confidence levels of the set of hypotheses Ω^L are updated, followed by updating the confidences of all nodes in the tree, Ω .




Fig. 3. Airflow model with Z (yellow crosses), E (purple cross with an arrow showing the predicted airflow direction, red is observed, blue is not), and resolve ambiguity locations (purple double sided arrows), for the scenario and robot trajectory shown in Fig. 6. Aqua arrows show the principal components of the clusters.

Algorithm 1 Determine conditional probability $p(E_k|\Omega_m^L)$

```

if ( $\Omega_m^L$  associated with  $E_k$ ) then
    high probability (0.8)
else
    if (corresponding evidence for  $\Omega_m^L$ ) then
        if ( $E_k$  associated with  $\Omega_m^L$ ) then
            high probability (0.8)
        else if (conflicting evidence for  $\Omega_m^L$ ) then
            very low prob (0.1)
        end if
    else
        if (flow with same dir in  $\Omega_m^L$  at location) then
            high prob (0.8)
        else if (there wasn't a flow in  $\Omega_m^L$  at location) then
            doesn't give any information (0.5)
        else if (conflicting flow in  $\Omega_m^L$  at location) then
            low prob (0.1)
        end if
    end if
end if

```

a) *Updating Ω^L (leaf nodes):* Therefore, for a given hypothesis Ω_m^L : given the prior probability $p(\Omega_m^L)$ determined through the airflow mapping, the conditional probabilities for each *observed* evidence $pr(E^{Obs}_j|\Omega_m^L)$ (for $(j = 1 \dots J)$) determined through expert knowledge encapsulated into Algorithm 1, the posterior probabilities $pr(\Omega_m^L|E^{Obs})$, can be evaluated using equation 1. This equation is an adaption of Bayes rule, for multiple evidences and multiple hypotheses. Subtleties of evidence are suppressed, and conditional independence among different evidences is assumed [16], avoiding the necessity to incorporate conditional probabilities of all possible combinations of evidences for all hypotheses.

b) *Updating Ω (all tree nodes):* The structure of the tree is retained, and the hypothesis probabilities are propagated up through the parents to the root nodes. The tree is traversed

$$p(\Omega_m^L | E^{Obs}) = \frac{p(E^{Obs}_1 | \Omega_m^L) \times p(E^{Obs}_2 | \Omega_m^L) \times \dots \times p(E^{Obs}_J | \Omega_m^L) \times p(\Omega_m^L)}{\sum_{q=1}^M p(E^{Obs}_1 | \Omega_q^L) \times p(E^{Obs}_2 | \Omega_q^L) \times \dots \times p(E_J | \Omega_q^L) \times p(\Omega_q^L)} \quad (1)$$




Fig. 4. Roma and environment.

post-order, (from leaves up to root nodes, breadth first). Each node's confidence is set to the sum of its children's, so that all node probabilities are updated and consistent.

5) *Select best hypothesis*: Although Ω^L have continually updated probabilities, the absolute value cannot be used. Leaves at a lower 'generation' will have a deceptively lower probability than leaves at a given higher 'generation', as they are a product of more 'decision' points, even though they may be on a branch that had a high probability at an earlier stage. Therefore, the entire tree structure is exploited. It is traversed in a depth/best first manner, where the 'best' child of a given node is determined as the child with highest confidence value.

B. Resolve ambiguity

In order to discriminate between two favoured hypotheses, it is necessary to determine where they differ, and move to those target locations to gather evidence. Therefore the sectors between hypotheses are compared, and those without a corresponding sector (similar morphology and location) are identified. Straight line segments comprising the sector boundary, that have differing flow direction in the other hypotheses are determined. The centres of these segments are used as a new set of target locations. The segments are shown in Fig. 3 for the trial displayed in Fig. 6.

V. EXPERIMENTAL SETUP

The robot used is named Roma, Fig. 4. It is equipped with a simple wind vane, only accurate to within 45° , and a Figaro Tin-Oxide chemical sensor. Roma has on-board low level control, and communicates with a PC where high level control and processing is carried out. The robot control is achieved with a combination of behaviour based and sense-map-plan-act systems. The reasoning layer issues commands to move to specific locations in the environment, and a Subsumption architecture (SubsuMe static library <http://sourceforge.net/projects/subsumelib>) is used to guide the robot to the location safely (using path planning and reactive avoid behaviours).

The environment covers approximately $2\text{m} \times 3\text{m}$ and is 30cm high. It is constructed with reconfigurable boxes, an enclosing plastic ceiling, and contains objects of various

shapes and sizes with height equal to the walls. Airflow is generated at an inlet at a rate of 0.5m/s. The odour is ethanol vapour, injected into the environment at a flow rate of 0.5ml/s

VI. RESULTS AND DISCUSSION

Several odour localisation trials were conducted for each of 5 scenarios (variations in room configurations), leading to a total of 11 trials. The scenarios displayed significantly different airflow topologies. These are illustrated for one scenario in detail (only 1 out of 2 trials is shown), Fig. 6, and in summary for the other four, Fig. 5.

All the figures show the environment from a birds eye perspective. In the schematic depictions, the target locations (sectors and flows) for the best hypothesis are shown (and labelled in Fig 5). Robot trajectories over multiple trials are shown with solid and dashed lines. The starting position of the robot is marked with an X. Numbers around the perimeter of the room and objects indicate odour source location candidates. Inlets/outlets are shown by solid black lines (inlets also have an arrow indicating air flow direction).

For 3 scenarios (Fig. 6 and Fig. 5a,d), the correct hypothesis was determined in every trial. For 2 scenarios (Fig. 5b,c), two equal best hypotheses were determined for all trials. In the case of (b), the inlet was correctly predicted to be at the left for both hypotheses, for all trials. A correct prediction could be made regardless of which hypothesis was chosen. In the case of (c) the wrong equal best hypothesis was chosen, and the trials were unsuccessful.

Despite some erroneous measurements and an imprecise wind vane, the method was shown to be effective in resolving the correct hypothesis, which in many cases was critical. For example, in the case illustrated in Fig. 6, the NaReM was given an incorrect physical map. The single hypothesis model would have predicted Hypothesis 3. As the odour source was placed at candidate 15, shown by X at the top left section of the environment, this would have led to the prediction that odour would be in high concentrations in the whole top half of the environment, rather than predominantly restricted to the small sector in the upper left, and would result in an incorrect odour source location prediction. However, the multiple hypotheses allowed correct modelling, and therefore led to successful localisation.

VII. CONCLUSIONS AND FUTURE WORK

This article has presented the use of a multiple 'Naïve Physics airflow model' hypothesis tree, and Bayesian inference, to select the best hypothesis as part of an odour localisation task. Experimental results over several trials on a variety of environment configurations, hence airflow topologies, demonstrate the ability to improve the robustness of the airflow modelling, which the odour localisation relies upon.

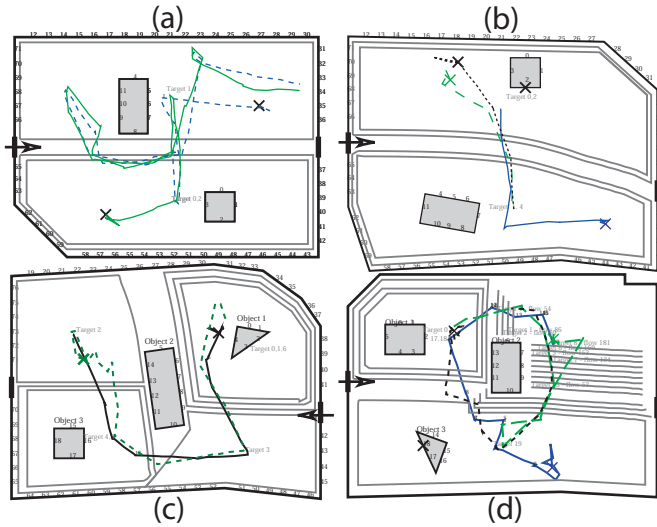


Fig. 5. Summary of experimental trials, showing the sectors of the best hypothesis, and the robot trajectories, for each of the scenarios.

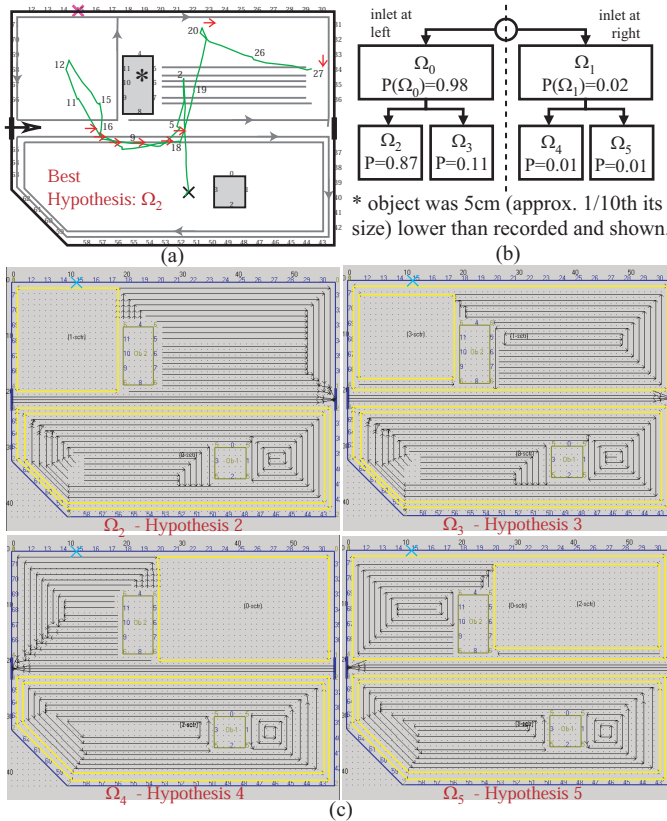


Fig. 6. Detailed results for one trial: a) schematic of best hypothesis with robot trajectory (labelled with timesteps) and arrows showing evidence location and predicted airflow direction (red=observed), b) hypothesis tree, and c) full airflow models of the leaf nodes (sector boundaries shown in yellow).

Initially, the ‘best’ hypothesis is Ω_5 . Roma takes a measurement at the closest target, and then for the next closest target (t_{0-2}). The updated confidences gives two equal best hypotheses - Ω_2 and Ω_3 . The best is Ω_3 , however, it does not contain any targets that have not been investigated. Therefore Ω_2 is set as the ‘current’ hypothesis. The robot then moves to all the remaining targets for Ω_3 (t_{3-12}). The best hypothesis remains unresolved, and Roma moves to the resolve ambiguity locations t_{13-27} .

This has important implications for other potential practical applications of Naïve Physics algorithms.

The most significant area for future development is the inclusion of other decision points. New hypotheses could be instantiated at additional decision boundaries, applying to other rules. Some examples were given in Section II-B. Resolving ambiguity could be extended with a wind vane capable of sensing air strength as well as direction.

ACKNOWLEDGMENT

Thank you to David Rawlinson for providing very useful code (<http://www.theCyberiad.net>) on which the particle filter was based, and to David Fernandez with whom the modified k-means algorithm was developed.

REFERENCES

- [1] H. Ishida, “Odor-source localization in the clean room by an autonomous mobile sensing system,” *Sensors & Actuators B*, vol. 33, no. 1-3, pp. 115–121, 1996.
- [2] F. W. Grasso, T. R. Consi, D. C. Mountain, and J. Atema, “Biomimetic robot lobster performs chemo-orientation in turbulence using a pair of spatially separated sensors: Progress and challenges,” *Robotics and Autonomous Systems*, vol. 30, pp. 115–131, 2000.
- [3] R. Russell, A. Bab-Hadiashar, R. Shepherd, and G. Wallace, “A comparison of reactive chemotaxis algorithms,” *Robotics and Autonomous Systems*, vol. 45, pp. 83–97, 2003.
- [4] H. Ishida, Y. Kagawa, T. Nakamoto, and T. Moriizumi, “Odor-source localization in clean room by autonomous mobile sensing system,” in *The 8th Int. Conf. on Solid-State Sensors and Actuators, and Eurosensors IX*, vol. 1, 1995, pp. 783–786.
- [5] A. T. Hayes, A. Marinoli, and R. M. Goodman, “Swarm robotic odor localization,” in *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, Wailea, 2001, pp. 1073–1078.
- [6] G. Kowadlo, D. Rawlinson, R. A. Russell, and R. Jarvis, “Bi-modal search using complementary sensing (olfaction/vision) for odour source localisation,” in *Proc. of the IEEE Int. Conf. on Robotics and Automation*, Orlando, 2006.
- [7] G. Kowadlo and R. Russell, “Naïve physics for effective odour localisation,” in *Australasian Conf. on Robotics and Automation*, Brisbane, 2003.
- [8] —, “To naïvely smell as no robot has smelt before,” in *Robotics, Automation and Mechatronics, 2004 IEEE Conf. on*, Singapore, 2004, pp. 898–903.
- [9] —, “Advanced airflow modelling using naïve physics for odour localisation,” in *Australasian Conf. on Robotics and Automation*, Sydney, 2005.
- [10] —, “Using naïve physics for odor localization in a cluttered indoor environment,” *Autonomous Robots*, vol. 20, no. 3, pp. 215–230, 2006.
- [11] D. Reid, “An algorithm for tracking multiple targets,” *Automatic Control, IEEE Transactions on*, vol. 24, no. 6, pp. 843–854, 1979.
- [12] W. Koch, “On bayesian mht for well-separated targets in densely cluttered environment,” *Record of the IEEE 1995 International Radar Conference*, pp. 323 – 8, 1995.
- [13] J. Huelsenbeck, F. Ronquist, R. Nielsen, and J. Bollback, “Bayesian inference of phylogeny and its impact on evolutionary biology,” *Science*, vol. 294, no. 5550, pp. 2310 – 2314, 2001.
- [14] I. J. Cox, J. M. Rehg, and S. Hingorani, “Bayesian multiple-hypothesis approach to edge grouping and contour segmentation,” *International Journal of Computer Vision*, vol. 11, no. 1, pp. 5 – 24, 1993.
- [15] I. J. Cox and J. J. Leonard, “Modeling a dynamic environment using a bayesian multiple hypothesis approach,” *Artificial Intelligence*, vol. 66, no. 2, pp. 311 – 344, 1994.
- [16] K.-C. Ng and B. Abramson, “Uncertainty management in expert systems,” *IEEE Expert*, vol. 5, no. 2, pp. 29–47, 1990.