Naïve Physics for Effective Odour Localisation

Gideon Kowadlo
Intelligent Robotics Research Centre
Monash University, AUSTRALIA
gkowadlo@ieee.org

R. Andrew Russell
Intelligent Robotics Research Centre
Monash University, AUSTRALIA
andy.russell@eng.monash.edu.au

Abstract

This paper describes current progress of a project that uses naïve physics to enable a robot to perform efficient odour localisation. Odour localisation is the problem of finding the source of an odour or other volatile chemical. Performing this effectively could lead to many humanitarian and other valuable applications. Current techniques utilise reactive control schemes requiring the robot to follow the plume along its entire length, which is slow and may be especially difficult in a cluttered environment. This research is concerned with creating a more 'intelligent' system to overcome these limitations. A map of the robot's environment was used, together with a naïve physics model of airflow to predict the pattern of air movement. The robot used the airflow pattern to reason about the probable location of the odour source. A prototype system was successful in a simplified cluttered environment, locating the source comparatively quickly. This demonstrates that naïve physics can be used for effective odour localisation, and has the potential to allow a robots operating in unstructured environments to reason about their surroundings. This paper presents details of the naïve physical model of airflow, reasoning system, experimental work, and results of practical odour source localisation experiments.

1. Introduction

Olfactory sensing is vital for many life forms on Earth, from bacteria to humans. Some species identify airborne odours, and locate the source to find food, nests and mates [Ishida, et al., 1999]. There are many valuable applications for a robot with the ability to locate the source of an odour/chemical plume. These include identifying the source of dangerous substances such as airborne radio-active or biological material, hazardous chemicals, gas and other pollutants in industrial and other settings; detecting unexploded mines and bombs; detecting illicit material at international ports; searching for survivors in earthquake damaged buildings, landslides or avalanches; detecting fire in its initial stages; and performing inter-robot communication.

Since the early 90's, many researchers have developed odour localising robots, employing a variety of methods to tackle the problem. Performance can be judged by the proportion of trials in which an odour source is located successfully, and the average time taken.

For most situations, the problem domain is an environment with a turbulent background fluid flow. Even indoors without an air source, small flows exist. Diffusion occurs extremely slowly, and the dominant means of odour dispersal is through carriage by the flow. The odour is released and creates an odour plume, which spreads out, mainly due to turbulence, as it travels downwind from the source.

For robots that are comparable in size to bacteria or for robots that burrow through the ground in search of chemical sources, fluid motion is governed by viscosity. The dominant means of odour dispersal is through diffusion, which creates smooth variations in concentration.

Research on odour localisation has been conducted predominantly for robots that operate in the former problem domain, within a background fluid flow. Most work has used reactive control schemes. The robot repeatedly responds to local sensing with small movements, to enable tracing of the plume all the way to the source. This approach is similar to the behaviour of some microbes, insects, and crustaceans.

Grasso [Grasso, et al., 2000] has built a biomimetic lobster that moves up a measured chemical concentration gradient. The lobster follows the plume-edge close to the source, where it is well defined and there is a significant gradient. Therefore it is successful if in close proximity to and facing the source.

Another method is to find the plume and then move upwind whilst remaining within the plume [Ishida, 1996], [Ishida, et al., 2002], [Russell, et al., 1995], [Hayes, et al., 2001].

Others have used techniques to obtain directional information using more than instantaneous gradient measurements. Active sensors of this type were built by Duckett [Duckett, et al., 2001], and by Ishida who created a sensor modelled on the silkworm moth [Ishida, et al., 1996], and a three dimensional odour compass [Ishida*, et al.*, 1999]. Passive sensors were used by Ishida [Ishida, 1996], [Ishida, 1994] and by Kazadi [Kazadi, et al., 2000].

To the authors' knowledge, there are only two studies that do not use reactive control schemes. Ishida used remote sensing and localisation by fitting a chemical distribution model to the sensor dynamics [Ishida, et al., 1997]. The robot must still move close to the source and makes repeated movements and calculations, which are time consuming. Russell has reported a robot that operates within simple interconnected tunnels. The method uses upwind searching and knowledge of gross fluid dynamics effects in this environment [Russell, et al., 2000], [Russell, 2001].

Direct comparison of the methods reviewed is difficult due to the variety of conditions. They were all successful at locating an odour source in their given environment in a similar time frame, from 5 to 15 minutes. Significantly, there are two main limitations of all the methods reviewed. Firstly, robots must follow the plume close to the source (in the case of [Ishida*, et al.*, 1997]), or to the source for identification. This may not be possible, and is time consuming. Secondly, none have dealt with the scenario of a typical indoor environment: a predominantly open area bound by walls and thinly populated by objects that affect airflow (although odour localisation in the presence of an obstacle has been looked at by Ishida [Ishida*, et al.*, 1999]). This is a severe limitation, as many of the applications of odour localisation will require the ability to operate in these conditions.

In order to overcome the limitations, we have taken the approach of implementing a 'sense-map-plan-act' control scheme, drawing on the type of intelligence that may be used by many higher-level animals, such as a human or dog, for similar sensor guided tasks.

This paper describes progress of the research. In initial work, a robot was able to quickly locate an odour source in a pseudo 2D indoor environment, with two objects, one of which was the source, and two openings (inlet and outlet). A 2D environment is used as a starting point to establish that the method as viable. The principles can be extended to 3D to make this applicable to real life applications.

The environment configurations were limited, and the robot required *a priori* knowledge of the direction of flow across the openings. Improvements of these aspects are proposed in this paper.

2. 'Intelligent' odour source localisation

Chemical sensors measure chemical concentration at the location of the sensor. They are not directional and do not usually form images. In order to infer the arrival direction of an odour, a windvane can be used to measure local airflow heading or more sophisticated odour sensors can be constructed to make the measurements directional. However, information of airflow direction is very local and cannot be used to predict a path all the way to the source. In order to achieve such behaviour, the robot must take more of the environment into account. It must obtain information about the room, such as the existence and positions of features (like doorways and objects), as well as the local information about odour concentration and wind direction. Then this information must be used effectively.

More of the environment can be taken into account by using complementary sensing. Complementary sensing is the use of multiple sensors, for a single purpose, that confer information about different aspects of the environment. It is not to be confused with sensor fusion, which is the use of multiple sensors to give information about the same aspect of the environment. In this research, complementary sensing comprises the use of 3D range data, for map building; combined with wind direction and odour sensing, for information about the odour plume.

The information about the environment is used effectively by creating a map of the airflow. Given complete information about airflow in every part of the robot's environment, odour plumes detected at the robot could be 'projected back' into the environment to identify probable locations for the source.

A conventional finite element model of fluid flow cannot be applied because many boundary conditions are unknown. Moreover, it would be extremely time consuming. In this project it is conjectured that a naïve physical model of fluid flow can be used to provide approximate information about patterns of airflow without requiring detailed knowledge of boundary conditions. A Reasoning System uses the knowledge of airflow from the airflow model to predict possible locations for the odour source and to propose movements of the robot that will provide additional information to improve the reliability of the predictions.

2.1 The naïve airflow model

Naïve physics is used to determine the airflow in the robot's environment. It is the use of common sense knowledge and physical intuition to model the environment, rather than the conventional approach of mathematical modelling. Take for example, the scenario of pouring liquid into a glass. We know what is going to occur to a level of detail that is sufficient for us to perform the operation effectively. We use naïve physics - common sense and physical intuition. Prediction using classical physics would involve solving a series of non-linear differential equations, requiring knowledge of all the boundary conditions. This would be time-consuming and difficult, if not impossible to achieve accurately. It provides more information than required, and the information is in a non-intuitive form that makes it difficult to extract simple properties of the behaviour.

The field of naïve physics began in the 80's with the naïve physics manifesto by Hayes [Hayes, 1979],[Hayes, 1985], which proposed a large-scale formalisation of commonsense knowledge. This was followed with other papers by AI researchers regarding theoretical and philosophical aspects of naïve physics. One group simulated some physical phenomena [Gardin and Meltzer, 1989], however, to the authors' knowledge, naïve physics has never been used for any practical applications, including robotics.

Hayes rejected attempts to solve practical problems using naïve physics because of the unrealistic and potentially misleading simplifications that are inevitably made. He reasoned that these attempts would ultimately not contribute to a system that has potential for *real* applications. We feel that a concurrent process of theoretical **and** practical work is required, each aspect guiding the other to produce progress in the right direction. In addition, we feel that it is possible to use the concepts of applying intuitive thinking to realise practical applications now. This is not intended solely as a means to an end, however we hope it is a study that can be built upon to further the field of naïve physics.

A general procedure, shown in Figure 1, has been developed for creating algorithms that use naïve physics to model an aspect of the environment. We have named the resulting algorithms **Na**ïve **Re**asoning **M**achines, NaReM's. In a sense, this is a supervised learning scheme. The system is trained to produce the same results as those observed (and used to derive the rules). Training has the potential to create robust algorithms that can generalise and therefore operate efficiently in unstructured environments.

In summary, NaReM's can model phenomena without knowledge of the mathematical relations, they require supervised learning, and have the ability to generalise. These features share some similarities with artificial neural networks (ANN's). However, ANN's operate on the signal, rather than symbolic level, and do not generalise well other than through interpolation.

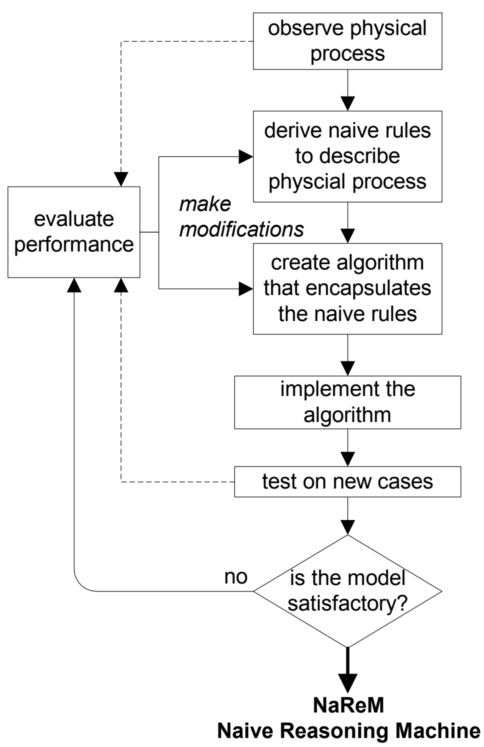

Figure 1. Creating a NaReM.

The stages shown in Figure 1 were carried out to create an airflow NaReM that was implemented and tested. An improved algorithm has also been developed and is currently being implemented. The stages of development for both are explained in the following sub-sections.

Initial airflow NaReM

<u>Observe phenomenon</u>

A large data set was required to cover as many effects as possible. To generate a large data set, simulations were conducted with a program called Flo++. Some key configurations were tested to verify simulation accuracy.

The simulations and subsequent algorithm development are limited to a constrained environment, the robot's world. It is defined as an area enclosed by walls and containing objects, both of which affect airflow that is introduced into and exhausted from the area via a number of openings.

As a starting point in the development of the method, the following constraints on the complexity of the system were made:

- The 'world' is a rectangular room with one inlet, one outlet, up to two objects, and an area in the order of 5m^2. The inlet and outlet are located on opposing walls.
- Airflow has the following characteristics:
    - The airflow inlet blows air perpendicular to the wall in which it was mounted.
    - The airflow source produces laminar (smooth, non-turbulent) airflow.
    - The airflow source produces airflows at a speed of approximately 0.5 m/s.

These airflow specifications are based on the fact that the airflow in an indoor environment is most likely to originate from an air-conditioning vent. Such vents usually have a grid with a depth that collimates the airflow, making it predominantly laminar and perpendicular to the wall in which it is mounted. In addition, they typically produce slow currents (in the order of 0.5 m/s).

The simulation results showed that sectors of airflow develop, and odour is restricted to the sector in which it is released. There is a small amount of mixing between sectors, most significantly due to turbulence. However, the dominant means of odour movement is by laminar airflow. Therefore, if an odour is detected at significant levels, and if a map of the room and the sectors of wind flow are available, it is possible to predict probable odour source locations. Therefore, sectors were identified in the simulation results and recorded as shown in Figure 2.

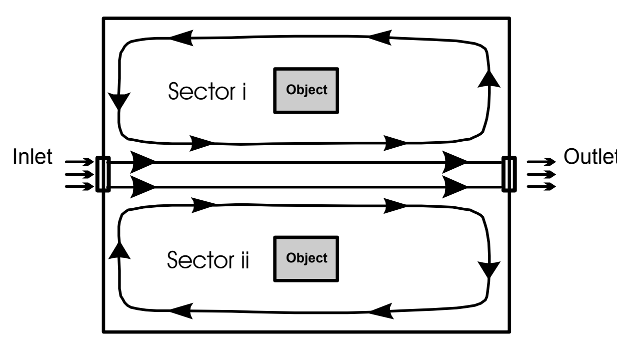

Figure 2. Features of an airflow simulation result.

Derive naïve rules that describe phenomenon

From the simulation results the following naïve rules were derived.

1. Air movement continues unless impeded.
2. If impeded, air stream bifurcates, and travels parallel to the obstructing object in both directions.
3. Airflows themselves act as obstacles to other airflows. However, air travels parallel to the obstructing airflow only in the same directions as this airflow.
4. Parallel airflows are conducted in a perpendicular direction (they spread out).
5. Air cannot be created or destroyed.
   - This rule is loosely applied. It does not need to hold locally in the end result. However, it must be observed when adding airflow arrows.
6. Real airflows are affected by solid boundaries - there is a reflection effect due to rules 2 and 4.
7. Opposing airflows are pseudo additive. Arrows of opposing directions cancel each other, and aligned arrows are retained.
8. Objects create airflow shadows, i.e. they block the air leaving an area of low pressure behind the side that faces the wind.

Create and refine NaReM that encapsulates naïve rules

These rules were formulated into an algorithm, which is divided into two sections: a first and second pass. It was designed to operate on a rectangular room, with an inlet and outlet that are positioned on opposite walls at x=0 and x=max. The room is divided into a grid with rows and columns in the y and x directions respectively. Each grid point has a **type** and a **direction**. The types can be one of:

WALL	-	Boundary to room.
VWALL	-	Wind, which acts like a virtual wall.
SOLID	-	An object.
EMPTY	-	nothing.

The direction, which applies only to type VWALL, can be horizontal (left or right), or vertical (up or down).

The algorithm is briefly described below to give an overview of its operation. The instructions are listed, followed by a number indicating the naïve rule on which they are based. The process is illustrated in Figure 3.

First Pass

- Draw an arrow starting from the inlet into the room (setting the airflow at each of the covered grid points to this direction). 1.
- If an obstacle is encountered, draw an airflow arrow, the size of the grid spacing, parallel to the obstacle in the 'scanning direction'. The 'scanning direction' is defined as perpendicularly away from this initial scan. This is the direction of subsequent scans, as explained further following this 'First Pass' description. 2.
  - Encountering an obstacle applies to the head and the tail of the airflow arrow. I.e. air must come from somewhere and go somewhere. Arrows are always added from head to tail. 5.
  - Obstacles can be a Wall, Virtual Wall (4) or Object.
- Airflow is not continued on the other side of the object. 8.

These steps are defined as a **scan**. The first pass begins with a **scan**, with the 'scan direction' set as from the inlet into the room. Then,

- Moving perpendicularly away from the initial scan, perform subsequent scans every grid spacing. Therefore, parallel scans are performed above and below the initial scan until the walls have been reached. 4.
- Each scan must initially extend the airflows that are perpendicular to the current scan (hence moving towards or away from the current scan). 1.

Second Pass

- Reflection from the walls: The same process is repeated, however the scan direction is reversed, and scans begin at the level of the walls and move towards the initial scan level of the input. 6.
- Superimpose the initial and reflection scans: If they are of equal length, then they cancel out, otherwise the stronger airflow swamps the weaker airflow. 7.

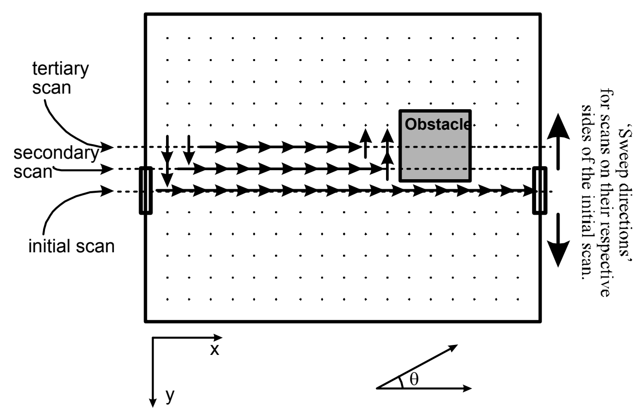

Figure 3. Airflow modelling.

The airflow-modelling algorithm works for a limited set of room configurations. The Reasoning System (see section 2.2) cannot operate effectively if the airflow map is inaccurate. It is possible to continue refining the current airflow algorithm, but ultimately it is restricted by its inherent non-symmetry. There is a distinction between airflows that are parallel to the scan direction and those that are perpendicular. This introduces difficulties with: having an outlet and inlet on perpendicular walls, generalising to non-rectangular rooms, and producing non-position specific air shadows from objects. Therefore a new more general algorithm is proposed, and is currently being implemented.

Proposed airflow NaReM

The proposed NaReM takes a very different approach at the Implement Algorithm stage. Airflow streams are represented by concurrent growing and if appropriate, snaking or bifurcating arrows. They move throughout the room, tracing out airflow directions. The room becomes divided into sectors. The process is then repeated within the sectors until there are no large empty spaces. Each repetition is an increase in 'level'. The higher the level, the lower the airflow velocity.

This algorithm has the potential to be very general, working for non-rectangular rooms with a variable number of obstacles placed in any position. Another advantage is that it can provide an indication of airflow velocity as well.

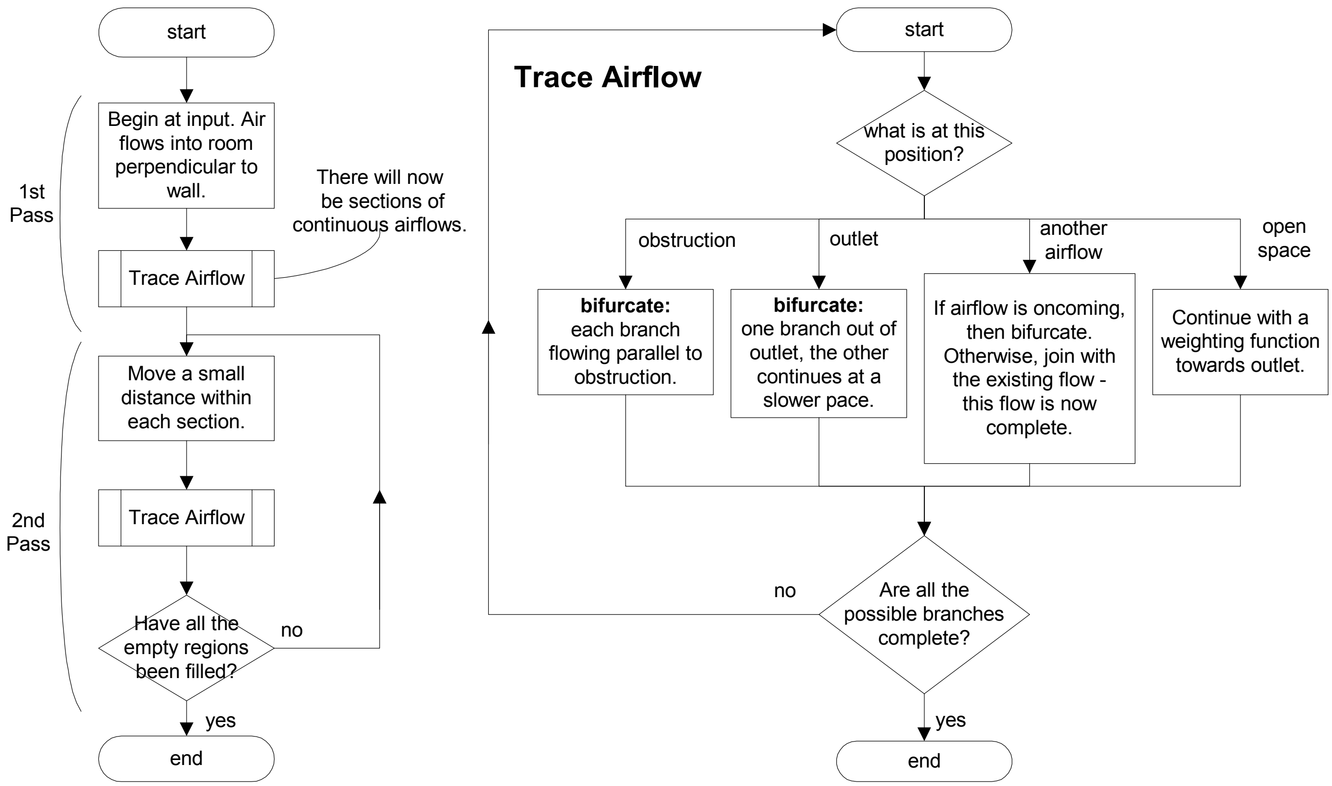

Figure 4. Airflow modelling algorithm proposal.

2.2 The Reasoning System for predicting the source location

The Reasoning System was designed to emulate the type of behaviour of an intelligent animal. It follows a 'Sense-Map-Plan-Act' control scheme that can be divided into four phases. The system is shown conceptually in Figure 5.

Phases two and three, the core of the system, have been implemented. Phase one is the construction of a map of the environment. Currently an *a priori* map is supplied to the robot. Phase two is the construction of a map of the airflow in the environment, using the airflow NaReM. Phase three is the prediction of the most probable odour source, using the airflow map as well as local sensing. This is performed by the Reasoning Algorithm, described in detail in the following section. Finally, phase four is the verification of the predicted odour source.

In future work, phase one will be accomplished with a laser rangefinder. The range data will be used to build a map. Phase four will involve the robot moving to and circumnavigating the predicted odour source. Verification will be affirmative if the concentration is very high on the downwind side, and low on the upwind side. In the event of a negative verification, the next most probable odour source would be investigated.

In the initial work, the *a priori* map included information about the direction of flow through the openings to the environment. Improvements in phase two that have now been implemented make this unnecessary. This is essential for the future implementation of phase one, autonomous map building, as this type of prior knowledge will be unavailable without an *a priori* map. All possible hypothesise are tested by creating airflow maps and comparing wind direction at the robot's location with the measured value. There may be more than one possible hypothesis, in which case the robot moves to the closest location where there is a conflicting predicted airflow direction for the two hypothesise, and the ambiguity is resolved.

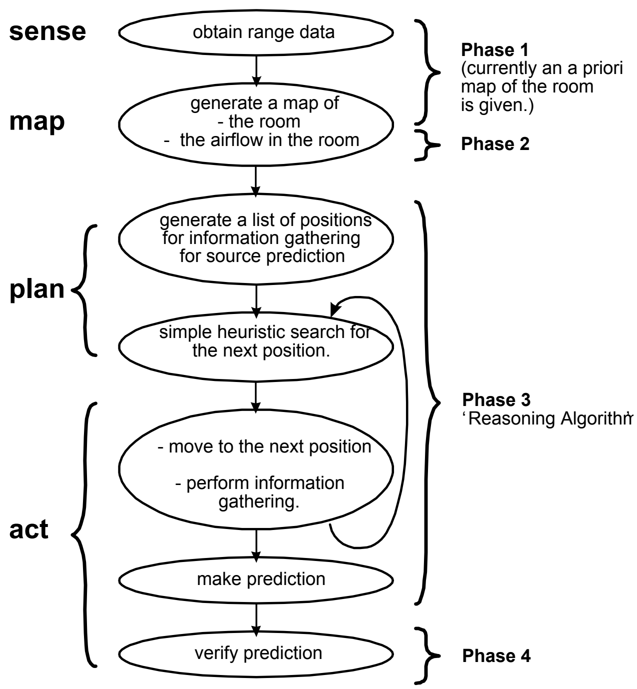

Figure 5. Reasoning System.

Reasoning Algorithm

The Reasoning Algorithm (RA) uses local sensing combined with the airflow map (produced by the airflow NaReM) to reason about the probable location of the odour source. The RA, summarised with the flowchart shown in Figure 6, is essentially a physical search through a list of key target positions, at which information is gathered to enable an odour source prediction.

It can be divided into three main sections, shown enclosed in dashed boxes in Figure 6. The first is to *create a list of target positions* by identifying the odour source candidates and tracing downstream from all sides of each candidate. Each trace must be investigated, but any position on the trace is suitable. This is a type of heuristic, as it narrows the search space of position to move towards, from the entire room, to a small list of locations.

Initially, and each time the robot moves a step, *information gathering*, comprising the second section, is performed.

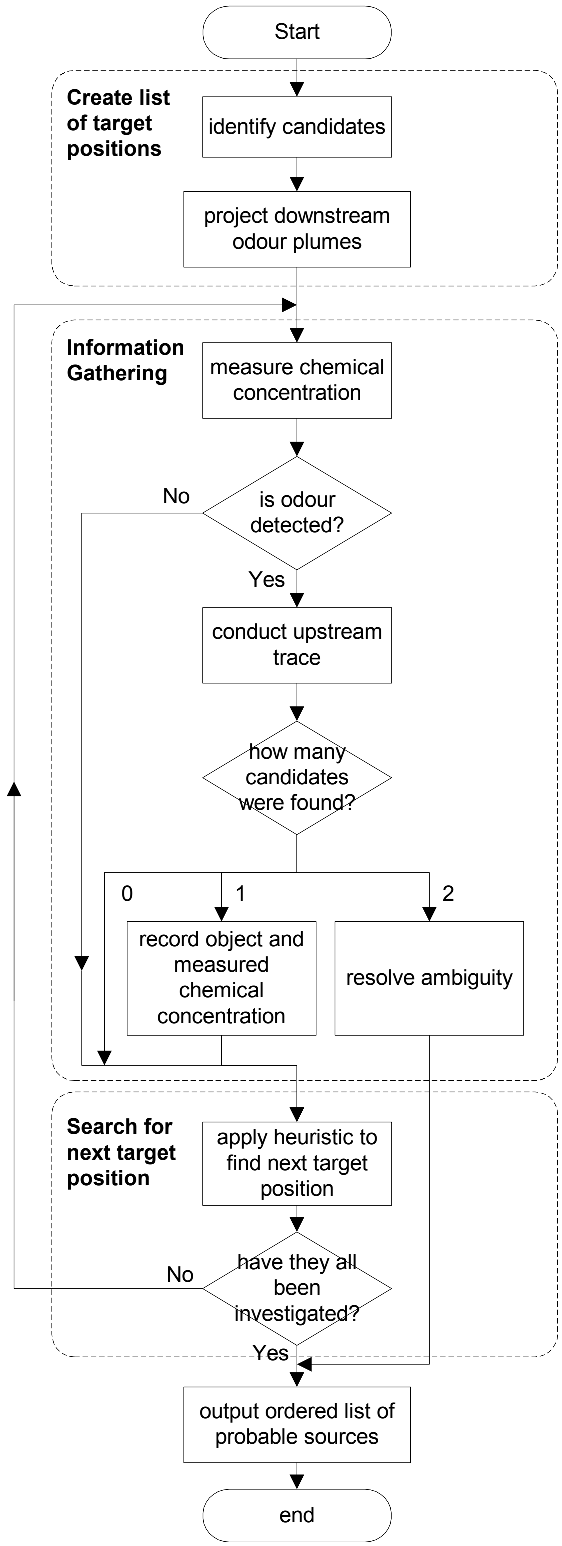

Figure 6. Flowchart of the Reasoning Algorithm.

After this, the RA must *search for the next target position*, comprising the third section. The objective is to minimise the total distance that must be travelled to explore all the traces. This is a variant of the 'travelling salesman' problem, with the added complexity that the target positions are not localised, being distributed linearly along a trace. It is solved effectively (though non-optimally) by applying the heuristic: find the closest position on a trace, out of all positions on all traces that have not yet been investigated. The robot then moves towards the target positions, reaching them after several steps, until every trace has been visited.

After investigating every downwind trace, a list of predicted sources is presented, ordered by the chemical concentration that was measured for that source. The first source is the primary prediction.

During *information gathering,* if odour is detected, a virtual trace *upstream* is conducted to identify which candidate(s) were responsible. If more than one candidate is identified, then the algorithm directs the robot to move to a unique position to resolve the ambiguity. In the case of two objects, and if one of the objects is downwind from the other, the position is the midpoint between them. If there are two objects but they are opposite each other with the wind passing between them equally, then it is not possible to determine which object is the odour source remotely, and one of them must be circumnavigated.

The 'downstream trace' is a quick method of determining target positions. However, once an odour is detected, it is necessary to ascertain if there are other possible candidates sourcing that position. The upwind trace does this by conducting a more complete search, taking the trace vicinity into account, looking for objects as well as other traces that are considered to be a source due to turbulent mixing. The upwind trace must also be used in the case of unpredicted odour detection (when not at a targeted position calculated as downwind from an object).

The robot is able to bypass objects with simple reactive control. The objects are ignored by the main reasoning algorithm, however, if during a move, a collision is detected, the robot will move backwards to clear the obstacle, rotate 90° in the direction of the target, and move forwards by 10cm.

3 Experimental Verification

3.1 Roma the robot - construction and interfacing

Roma, a mobile two-wheeled robot, was constructed and used to test the odour localisation method. It has a circular 24cm diameter base and is approximately 15cm high. Roma can be commanded to move forwards, backwards, to turn on the spot both clockwise and anticlockwise, and to report sensor readings. Roma is shown in Figure 7.

Roma consists of a chassis, robot drive system, sensors, support circuitry and a microcontroller. It was connected to a PC and a power source via a tether. The Reasoning System is run on a PC, which sends commands to the microcontroller via an RS232 link. A program running on the microcontroller performs low level robot control. The robot system is shown in Figure 8.

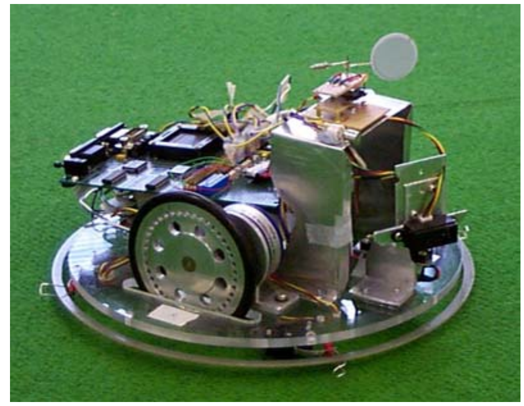

Figure 7. Photograph of Roma.

A graphical user interface is run on the PC. It was constructed using QT from Trolltech to facilitate development and testing of the robot.

The robot drive system comprises a passive castor and two wheels that are mounted along the diameter. Each wheel has an independent gearbox and dc motor for differential steering. With this arrangement, Roma is able to rotate on the spot about its centre.

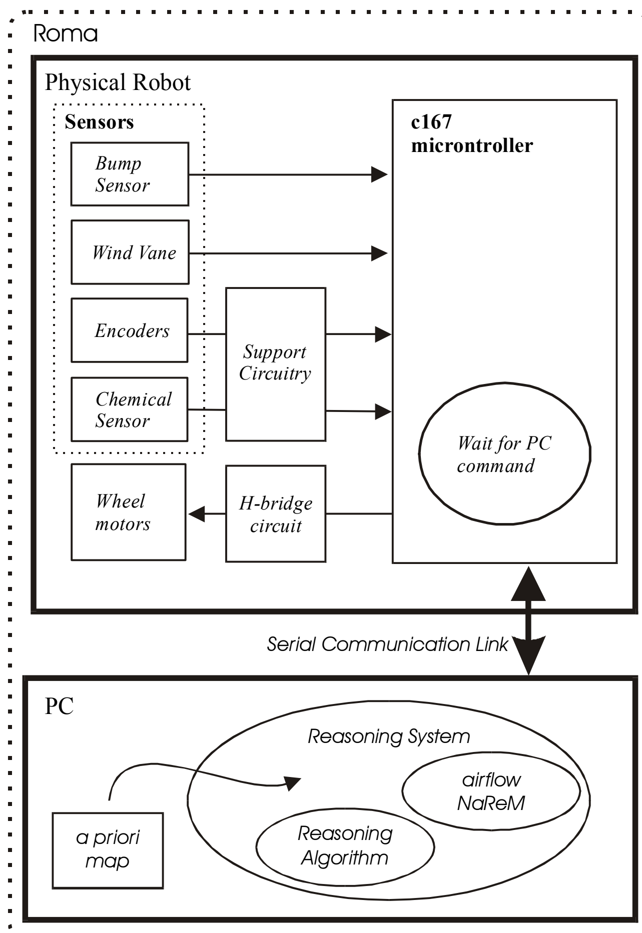

Figure 8. Robot systems.

The sensors include a wind vane, chemical sensor, bump sensor, and wheel encoders. The wind vane, which is mounted on top of the robot, can be seen in Figure 7. It measures the airflow direction to within one of eight 22.5º ranges. Odour is sensed with a Tin-Oxide chemical sensor. Their conductivity is dependant on the concentration of reducing gasses such as methanol, ethanol and ammonia. The response time is slow, in the order of seconds, and they are sensitive to changes in temperature and humidity. A bump sensor ring surrounding the robot detects collisions. Optical encoder sensors are used for feedback of wheel rotation. The angular resolution of turning is 4.7°, and linear resolution of movement is 7.5mm.

The microcontroller is an Infineon C167. It performs 'Robot Control': controlling the Robot's movement, acquiring data from the sensors, receiving commands from the Reasoning System (running on the PC), and responding to those commands by either moving or sending data back to the PC. The control system for movement is essentially a feedback proportional controller.

3.2 Roma's world

The odour localisation method was developed and tested in a pseudo 2D room. Two-dimensionality was simulated by creating a room with a height that was small compared to the other dimensions and with objects and airflow generator that extended from the floor to the ceiling.

The room, shown in Figure 9, had a floor area of 2820 x 1900 mm. A single layer of re-configurable boxes 270 x 270 x 470 mm were used to build the walls and clear plastic, attached to a metal frame, formed the ceiling. Two boxes, with a cross section of 260 x 260 mm, were placed in the room as potential odour sources. They had a height equal to the ceiling, and were large enough to disturb airflow. One of the objects was an odour source, which the robot was required to locate. A laminar air source introduced airflow into the room at a flow rate of up to 0.5 m/s.

(a) (b)

Figure 9. Photograph of the room and laminar air source.

The laminar air source, shown in Figure 9, consists of six DC fans that blow air through a collimator made up of an array of straws, 205mm in length and 4.5mm in diameter.

The odour source was created by injecting ethanol vapour from one side of the box at a flow rate of 0.5 ml/s. The vapour was generated by bubbling air through a flask of methylated spirits, causing it to become fully saturated.

3.3 System evaluation

With the Reasoning System set up to use the first NaReM, Roma was tested with ten trials. Three room configurations were used with one inlet, one outlet and two boxes (one of which was an odour source). The inlet and outlet positions were fixed, and the box positions as well as initial robot position were varied. For each trial, the odour source was switched on 10 seconds prior to activation of the robot to allow the chemical to circulate around the room and reach a steady state pattern. In order to compensate for humidity and temperature variations, the sensor threshold for odour detection was adjusted manually.

A selection of the results (with two of the three room configurations) are shown in Figure 10 (a,b). Each figure shows the robot's movements for two trials within that configuration. The inlet/outlets are shown as narrow rectangles with an arrow to indicate airflow direction, and the predicted downwind odour traces from the odour source candidates are shown with a dotted/dashed line (see legend). For each trial, the robot's initial position is indicated with a large circle, the final position with a cross, and the path taken by a line, with each step delineated by a circle. The circle is bold where a collision has occurred, and the subsequent movements numbered chronologically. The stages of the robot's behaviour are indicated with line style and symbols, which are explained in the legends.

Roma successfully identified the correct object as being the odour source in all of the ten trials. Localisation was achieved expediently, with a typical time of 30 – 50 seconds compared to 6 to 15 minutes, reported in the literature, for other single robot methods operating in the same or a smaller area. The success of the experiments is encouraging. It shows that this new approach to odour localisation is feasible and can be effective.

These tests were conducted with the initial NaReM. As explained in section 2.1, this is effective for a limited set of room configurations, restricting the utility of the Reasoning Algorithm. An improved airflow NaReM is therefore proposed in section 2.1. It is currently being implemented.

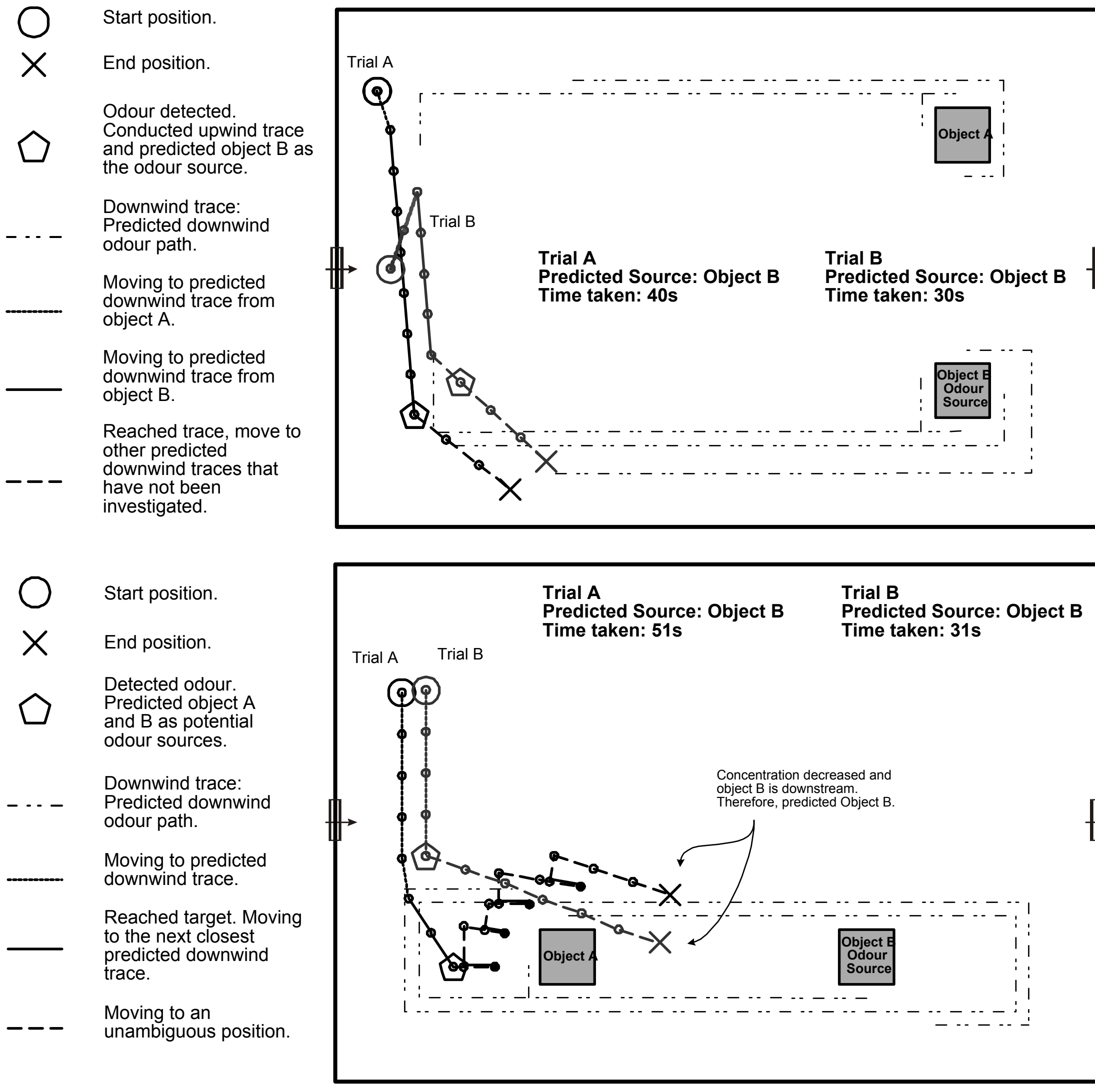

a)

b)

Figure 10. Experimental Results

Conclusion

Robots with the ability to locate an odour source have many useful and humanitarian applications. At present, odour localising robots use reactive control schemes to follow a plume along its entire length. This is slow and may not be possible in an indoor environment with obstacles, such as the situations that would be encountered during search and rescue operations. These issues have been addressed with a robot named Roma that uses a 'sense-map-plan-act' control scheme. It aims to emulate the type of odour source locating behaviour that may be employed by intelligent animals. Roma uses a novel **Na**ïve **Re**asoning **M**achine, NaReM, which uses naïve physics, to model the airflow in its environment. A reasoning algorithm then uses knowledge of the airflow to forecast odour paths, move to selected positions, gather information and make a prediction of the most probable odour source. The results show that this method can be effective in a simple, known, cluttered indoor environment. The robot does not need to travel all the way to the source, resulting in robust and comparatively fast performance. There is great potential for this approach to lead to many valuable applications.

Furthermore, through implementing the odour localisation approach, we have developed a method for using naïve physics for practical applications This has enabled, to the author's knowledge, the first successful application of a naïve physics model to a real robotics task, and demonstrates the potential of Naïve Reasoning Machines. If robots are to operate effectively in unstructured environments, they will need to be able to predict the consequences of their actions as well as changes in their environment. In order to do this, they will require a commonsense reasoning ability. Development in this area holds the key to advances for robust useful robots. This research involves the application of naïve physics to the specific task of modelling airflow and through this, odour localisation. However, as the initial success in this application has shown, there is great potential for more general use of naïve physics within the field of artificial intelligence, and more specifically, robotics.

Future work will include two main areas of improvement: the construction of the proposed airflow NaReM, and the implementation of autonomous map building using a rangefinder. Another major addition will be a predicted odour source verification stage, whereby the robot circumnavigates the most probable odour source. These improvements will enable the robot to move into unknown unstructured environments, and confirm predicted odour sources.

Although the project is at an early stage, our novel method has offered a more effective technique for odour localisation, and has demonstrated the potential of the approach for many valuable applications. It has also demonstrated the potential of naïve physics for robots operating in unstructured environments.

Acknowledgements

This project has been supported by grants from the Monash University Harold Armstrong Memorial Fund and Monash Engineering Faculty Grant Scheme. The authors wish to thank Tony Brosinsky and Maurice Gay for their assistance with the mechanical construction.

References

[Duckett, *et al.*, 2001] T. Duckett, M. Axelsson and A. Saffiotti. Learning to locate an odour source with a mobile robot. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA '01)*, 2001.

[Gardin and Meltzer, 1989] F. Gardin and B. Meltzer. Analogical representations of naive physics. *Artificial Intelligence*, 38(2):139-159, 1989.

[Grasso, *et al.*, 2000] 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*, 30:115-131, 2000.

[Hayes, *et al.*, 2001] A. T. Hayes, A. Marinoli and R. M. Goodman. Swarm robotic odor localization. In *Proceedings of the IEEE Conference on Intelligent Robots and Systems (IROS '01)*, pages 1073-1078, 2001.

[Hayes, 1979] P. Hayes, J. The naive physics manifesto. in D. Michie ed. *Expert systems in the microelectronic age*, Edinburgh University Press, 1979, pages.

[Hayes, 1985] P. Hayes, J. The second naive physics manifesto. in J. Hobbs and R. Moore eds. *Formal theories of the commonsense world*, Ablex Publishing Corporation, Norwood, NJ., 1985, pages 1-36.

[Ishida, 1994] H. Ishida. Study of autonomous mobile sensing system for localization of odor source using gas sensors and anemometric sensors. *Sensors & Actuators A*, 45(2):153-157, November 1994.

[Ishida, 1996] H. Ishida. Odor-source localization in the clean room by an autonomous mobile sensing system. *Sensors & Actuators B*, 33(1-3):115-121, July 1996.

[Ishida, *et al.*, 1996] H. Ishida, K. Hayashi, M. Takusaki, T. Nakamoto, T. Moriizumi and R. Kanzaki. Odour-source localization system mimicking behaviour of silkworm moth. *Sensors and Actuators A*, 51:225-230, 1996.

[Ishida, *et al.*, 1999] H. Ishida, A. Kobayashi, T. Nakamoto and T. Moriizumi. Three-dimensional odor compass. *IEEE Transactions on Robotics and Automation*, 15(2):251-257, April 1999.

[Ishida, *et al.*, 1997] H. Ishida, T. Nakamoto and T. Moriizumi. Remote sensing and localization of gas/odor source and distribution using mobile sensing system. In *International Conference on Solid-State Sensors and Actuators*, pages 559-562, 1997.

[Ishida, *et al.*, 2002] H. Ishida, G. Nakayama, T. Nakamoto and T. Moriizumi. Controlling a gas/odor plume-tracking robot based on transient responses of gas sensors. In *Proceedings of IEEE Sensors, 2002.*, volume 2, pages 1665 -1670, 2002.

[Kazadi, *et al.*, 2000] S. Kazadi, R. Goodman, D. Tsikata and H. Lin. An autonomous water vapor plume tracking robot using passive resistive polymer sensors. *Autonomous Robots*, 9(2):175-188, 2000.

[Russell, 2001] R. A. Russell. Tracking chemical plumes in constrained environments. *Robotica*, 19(4):451-458, July 2001.

[Russell, *et al.*, 2000] R. A. Russell, L. Kleeman and S. Kennedy. Using volatile chemical to help locate targets in complex environments. In *Proceedings of the Australian Conference on Robotics and Automation*, pages 87-91, 2000.

[Russell, *et al.*, 1995] R. A. Russell, D. Thiel, R. Deveza and A. Mackay-Sim. A robotic system to locate hazardous chemical leaks. In *Proceedings of 1995 IEEE International Conference on Robotics and Automation*, volume 1, pages 556 -561, 1995.