

Advanced Airflow Modelling Using Naive Physics for Odour Localisation

Gideon Kowadlo and R.Andrew.Russell

Intelligent Robotics Research Centre

Monash University, Australia

gkowadlo@ieee.org

andy.russell@eng.monash.edu.au

Abstract

To date, robotics has had limited success at operating in unstructured environments. Part of the problem is the lack of commonsense reasoning. One area of commonsense reasoning is Naive Physics, the practice of using intuitive rules to reason about the physical environment. Researchers have explored relevant philosophical issues, attempted to develop logical formalisations, and developed systems to ‘understand’ and learn simple physical/spatial qualities of physical objects. This paper reports the implementation of an algorithm that uses Naive Physics to model airflow, and has been used on a practical robot for the task of odour localisation. This is an example of a Naive Reasoning Machine, an algorithm encapsulating naive rules, which represents a broader and more practical approach to naive physics.

1 Introduction

Despite great advance in artificial intelligence, “the AI community has been unable to create successful, autonomous robots able to operate in complex, unconstrained environments” [Boxer, 2001]. This has been attributed to an inability to emulate ‘commonsense reasoning’. A division of commonsense reasoning is Naive Physics, or the use of intuitive rules to understand and reason about the physical environment.

Naive physics was first brought into prominence by Hayes with his Naive Physics Manifesto [Hayes, 1979]. His passionate manifesto called on the community to develop a large scale formalisation of commonsense rules, that could be used for artificial intelligence capable of dealing with unstructured environments rather than ‘toy world’ scenarios. This triggered many research publications within the field of artificial intelligence, but also in other fields including geography [Egenhofer and Mark, 1995] and psychology [Watt, 1995]. The

bulk of attention from the AI field formed a divergent branch termed ‘qualitative physics’ [Dohnal, 1991; Kuipers, 1994], which strives to perform mathematical calculations more qualitatively, dealing with ranges rather than accurate quantitative information. Others continued to deal with the philosophical aspects [Smith and Casati, 1994] and fundamental questions regarding the form that the formalisations should take [Davis, 1997]. Gardin [1989] used analogical representations to model the behaviour of strings and liquids. In addition, there have been efforts in the area of logical formalisation with answers to the ‘egg cracking problem’ challenge put forward by Davis, to encapsulate the commonsense knowledge of the outcome of cracking an egg over a bowl [Lifschitz, 1998; Morgenstern, 2001; Shanahan, 2004]. Forbus [1980] has developed a system to reason about motion qualitatively, and recently, [Boxer, 2001] used a Bayesian network to learn qualitative spatial representation by visual observation.

Naive physics is often defined as “the ability to learn and reason about the physical behaviour of objects [Boxer, 2001]”. We have taken a broader definition, including the ability to reason about the environment in general. We see it as not just enabling AI systems to possess a general commonsense reasoning capability, but to use a set of rules to construct an algorithm, or a Naive Reasoning Machine (NaReM) for modelling an aspect of the environment (for information on the methodology used to develop NaReMs, see [Kowadlo and Russell, 2003; 2004]). This has advantages over conventional mathematical approaches as it does not require the use of complex mathematical solutions. These are often difficult to set up, cannot be solved without precise boundary conditions, and are time consuming. In addition, conventional mathematics provides detailed quantitative results that are difficult to use. NaReM’s, in part by their use of analogical representations [Gardin and Meltzer, 1989], produce results at a level of abstraction that can be readily exploited by a higher level reasoning process.

We have developed a NaReM to model the airflow in

an enclosed environment as part of a robotic odour localisation system. The initial implementation was successful for only very simple environments [Kowadlo and Russell, 2003], mostly due to the indirect application of the rules on non-analogical data structures, as well as a small number of rules.

An improved implementation was developed to address these issues, which was successfully used to ascertain the general location of odour sources in a variety of configurations of an environment consisting of an enclosed area (approximately 2m x 3m) with up to three objects and with an inlet and outlet. An overview of the system and results was presented in [Kowadlo and Russell, 2004], with a brief explanation of the NaReM. This paper explains the NaReM in detail, giving results of this component, and the future work. It begins with background work to put it into context, including an explanation of odour localisation.

2 Background Work

2.1 Odour Localisation

Odour localisation is the task of finding the source of an odour. It has received growing attention over the past 15 years. Much of the work has been inspired by the animal world, which contains many examples of organisms that perform this task to find mates or food. Odour localising robots could be used for a diverse range of applications and environments, including finding people in disaster zones such as buildings that have been damaged by explosives or earthquakes, finding gas leaks in industrial contexts, locating petrol leaks underground, or to gain a greater understanding of the behaviour of animals that rely on odour localisation.

The diversity of applications is reflected in the research publications. Russell has developed robots to locate chemicals in corridors [Russell, 2001] and underground [Russell, 2004], Grasso [2000] has developed a biomimetic lobster for underwater operation and the majority of work has focused on above ground operation in largely open areas. In this context, odour dispersal is dominated by carriage as a result of the background fluid flow, which is turbulent. The robots developed have predominantly used reactive control systems that follow the plume to the source. There has been a progression from robots that use plume following with only chemical sensors [Rozas *et al.*, 1991; Genovese *et al.*, 1992], to those that find the plume and then follow it upwind using chemical and airflow sensors [Ishida *et al.*, 1996; Russell, 1998; Marques *et al.*, 2003], and recently more sophisticated behaviours have been developed using vision as well [Ishida *et al.*, 2004; Loutfi *et al.*, 2004].

We have addressed odour localisation in enclosed environments that may be encountered in mine shafts, caves,

damaged buildings and industry. In this environment the airflows circulate forming sectors that may have very little intermixing. Chemical released into a sector will spread throughout that sector, making it difficult and slow to employ the conventional reactive plume following behaviour. Instead, a more sophisticated approach is required.

2.2 Odour Localisation System

Our solution [Kowadlo and Russell, 2003; 2004] involves initially modelling the airflow in the robot's surroundings. The prediction is then used to determine the physical locations that will maximise the utility of chemical concentration measurements. The robot moves to those positions, taking readings, and uses the information to predict the location of the odour source. The airflow is computed using a NaReM, which is the topic of this paper.

The robot is provided with an *a priori* map of the enclosed environment, which is no larger than 2m x 3m, with a height of 50cm, contains one inlet and one outlet, and up to three boxes of the same height as the ceiling.

3 Airflow NaReM

3.1 Naive Rules

A set of naive rules to describe the airflow were derived using a large data set of airflow patterns within the environment. These patterns were derived by running a range of simulations (using fluid dynamics software called Flo++), and key 'real world' experiments to: a) verify simulation accuracy, and b) ascertain the limitations of the simulations. The most important feature of the naive rules is that they must be capable of modelling the airflow sectors. They are described below: Note: For all rules, flows are only initiated given adequate empty space.

1. Air flows out of inlets and into outlets.
2. Flows have inertia, an impetus to continue moving in the same direction at the same strength, referred to as momentum direction, and momentum magnitude.
3. Flow continues unless impeded by an obstacle.
 - (a) An obstacle can consist of an airflow, object, or wall.
4. In open space, airflow is attracted to the outlet(s).
 - (a) The attraction causes the flow to be directed towards the outlet.
 - (b) The attraction decreases exponentially with increasing distance from the outlet.
5. Airflows spread out i.e. new flows form, given open space, on either side of existing flows.

- (a) These new flows flow parallel to the primary flow.
 - (b) They have a diminished strength.
6. If a flow encounters a solid obstacle (object or wall):
- (a) If the angle formed between the flow and the obstacle is small, the flow continues parallel to the surface of the obstacle.
 - (b) Else, the flow bifurcates into two flows that continue parallel to the surface of the obstacle in both directions. These flows have diminished strength
7. When a flow is following a solid obstacle's surface:
- (a) If the flow undergoes a large enough change in direction, then it becomes disconnected from the object, and propagates independently.
 - (b) As flows flow parallel to a solid obstacle's surface, the momentum direction changes incrementally to be closer to parallel with the surface it is following. This affects the flow propagation following disconnection from the obstacle.
8. If the flow encounters another flow:
- (a) If the flows are oncoming:
 - i. If the flows have the same strength, then both flows bifurcate. Where a flow branches, the branches have diminished strengths.
 - ii. Else, the stronger flow 'wins', overpowering or removing the weaker flow.
 - (b) Else, if the intersected flow has significantly greater strength:
 - i. Then branch, following parallel to the intersected flow.
 - (c) Else, the flow merges with the intersected flow. It thereby becomes a 'source' to that flow, and is considered a parent i.e. if it was carrying odour, the parent carries that odour to the child.
9. If a flow is flowing parallel to a wall and it passes an outlet, it bifurcates, with one branch flowing to the outlet, and the other continuing on the same path with diminished strength.

3.2 Encapsulation Into an Algorithm

Representation of Data

Flows and Flow Points The flows are represented as linked lists of flow data points FDP's (see Figure 1). A Flow Data Point (FDP) consists of a coordinate, momentum magnitude, momentum direction, and flow direction. The FDP is continually moved as the flow propagates. When the flow undergoes a deflection, then a new FDP is initiated. Therefore, a flow is comprised of

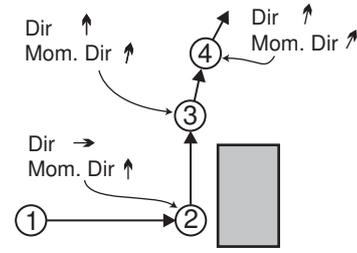


Figure 1: Illustration of flow progression

straight-line segments, which reduces memory requirements, and is computationally efficient. The flow direction describes the previous segment and the momentum is used to determine the next FDP in order to propagate the flow, explained further in section 3.2 'Incrementing Flows'.

Storing Flows

References to the flows are stored in two separate containers: a hybrid data structure consisting of a doubly linked list and graph; as well as a grid. The graph allows detection of sectors at run time as the nodes are added (see section 3.2 'Creating Sectors'). The doubly linked list allows a convenient way to exhaustively cycle through every flow. The grid allows random access to flows that occur in a specific physical location. This significantly speeds up searching for a possible collision between two flows, or a flow and a physical object.

Illustration of Flow Propagation

Figure 1 illustrates several rules and the process of incrementing the flow. The FD points are numbered chronologically from 1 to 4 (FD1-FD4). FD1 advances in a straight line until impeded by an obstacle (Rule 3), at which point the flow bifurcates, and a new FDP is added (FD2). The obstacle surface acts as a 'master', guiding both branches (Rule 6b). Concerning the upper branch: FD3 advances parallel to the object, until it reaches a point where the required angle to remain parallel is too great (Rule 7a), and it becomes 'disconnected' from the master. By this stage, the momentum direction is almost pointing parallel to the object (Rule 7b). Without a master, and not in open space, the momentum direction is used to advance the flow point (Rule 2). At FD4, there is open space, and the momentum is affected by the proximity of an outlet (Rule 4), which causes a deflection in the flow.

Operating on the Data Structures

The rules have been formulated into an algorithm, which is expressed in the flowcharts below. The first flowchart, Figure 2, shows the algorithm at the highest level. Thereafter, the subsections are expanded upon in further flowcharts.

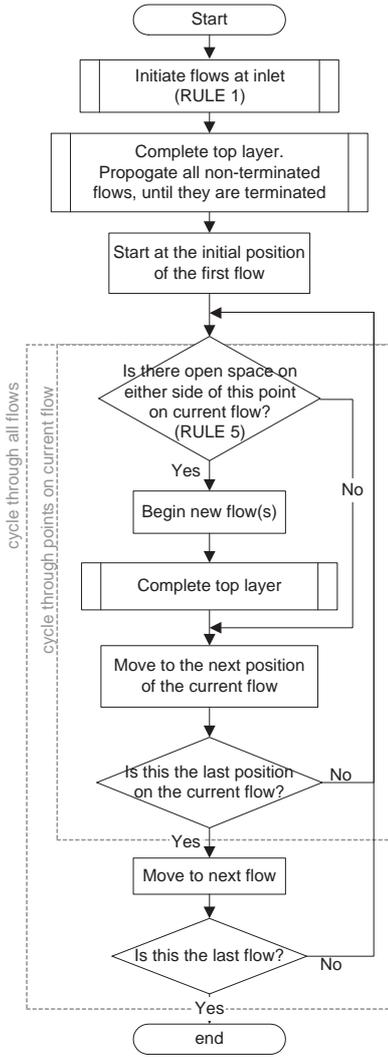


Figure 2: NaReM top layer

Complete Top Layer Cycle through all the active flows, incrementing them appropriately according to the rules, until they have been terminated. Each time a layer is completed, it comprises a new ‘layer’.

Initiate Flows At Inlet Two flows are begun at each extreme point of the inlet, traveling perpendicular to the inlet into the room. If the inlet is large enough, more than one flow will be initiated between these two points. This models air coming from the whole inlet, but also allows discrete sectors to be created if for example there is a bifurcation as a result of a solid object in front of the inlet.

Incrementing of flows, governed by the naive rules, is achieved using the algorithm expressed in the flow charts shown in Figures 3 and 4.

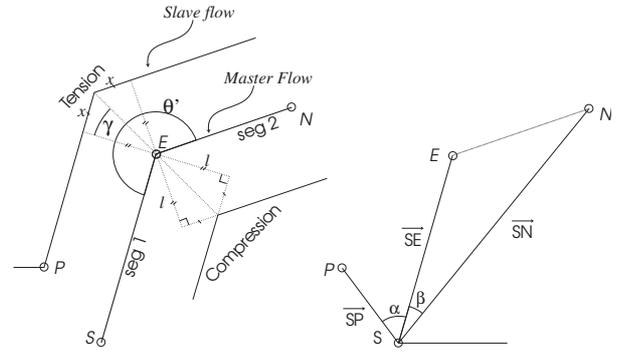


Figure 5: Slave and Master Flows: P = current flow point, S = start of segment, E = end, N = next

Incrementing Flows When air flows near an obstacle or another airflow, it can become attached. After attachment, it is considered to be a ‘slave’, with the obstruction the ‘master’. The flow (slave) is then incremented parallel to the master, and its momentum angle slowly changes towards the angle of the master.

A master point is incremented along the master flow, and the slave is incremented accordingly. At a corner, the flow must continue for a greater distance if on the outside (referred to as tension), than on the inside (referred to as compression).

The determination of tension or compression is illustrated in Figure 5. It is achieved by constructing and inspecting vectors SP, SE and SN. The angles α and β are calculated: α is the angle of SP with respect to SE and β the angle of SN with respect to SE. SE is treated as the origin, and the sign of the angles calculated accordingly. If α and β have the same sign, then Point P is in compression, otherwise it is in tension. The length of the flow is extended or reduced by x calculated in equation 1.

$$x = l \cdot \tan \gamma \quad (1)$$

$$\gamma = \frac{\theta'}{2} - \frac{\pi}{2} \quad (2)$$

Where l is the perpendicular distance maintained between flows.

If there is no master, then the flow can be propagated by first or second order (depending upon the stage of the algorithm). In the simple case of first order, the flow is propagated as per equation 3 and 4.

$$x_{i+1} = x_i + k |M| \cos(\angle M) \quad (3)$$

$$y_{i+1} = y_i + k |M| \sin(\angle M) \quad (4)$$

where M = momentum and k is a constant.

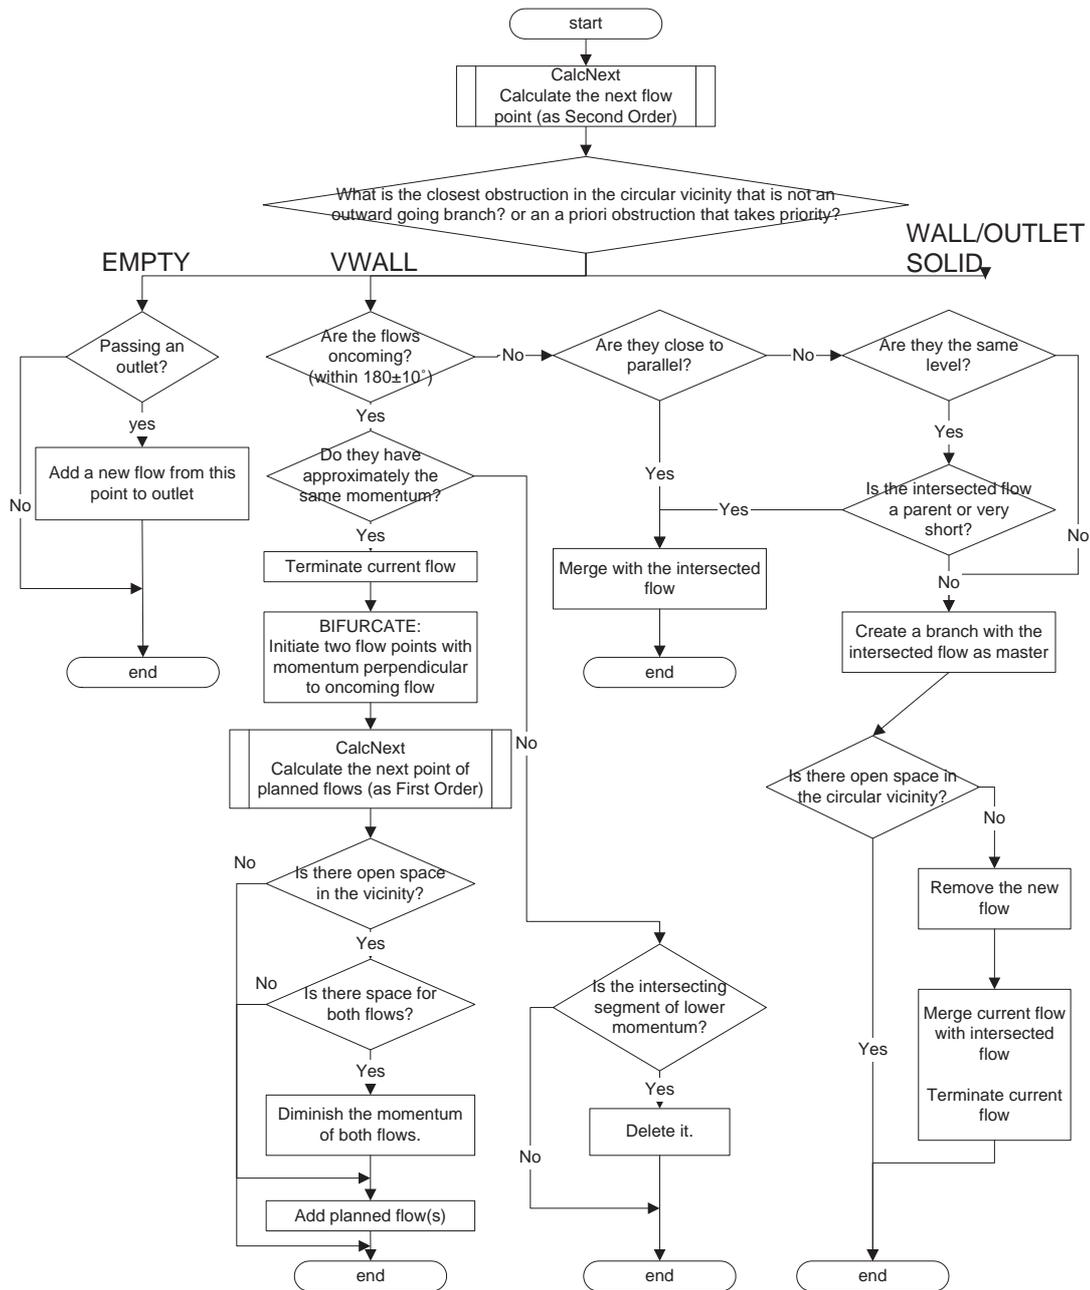


Figure 3: Increment flow: part 1

If second order, then the outlet and open spaces are taken into consideration. When in open space, the outlet attracts the flow. The attraction decreases exponentially with greater distance from the outlet.

The key steps of the process are shown in the flowchart of Figure 6.

Creating Sectors

Looking at a small section of the indoor environment, the airflow pattern becomes established as shown in Figure 7. The numbers indicate flow labels.

If a flow collides with another flow, the secondary flow is considered to be a child and is added as such in the flow graph. The physical basis is that the parent flow becomes a 'source' of the child, in terms of air and airborne chemical. Each node has one or more lineages, where a lineage is a list of ancestors that lead to this node. The lineages are used to detect when the addition of a new node creates a closed cycle, representing a sector. A lineage is comprised of units, where each unit consists of a flow and a coordinate. This is necessary, as not all

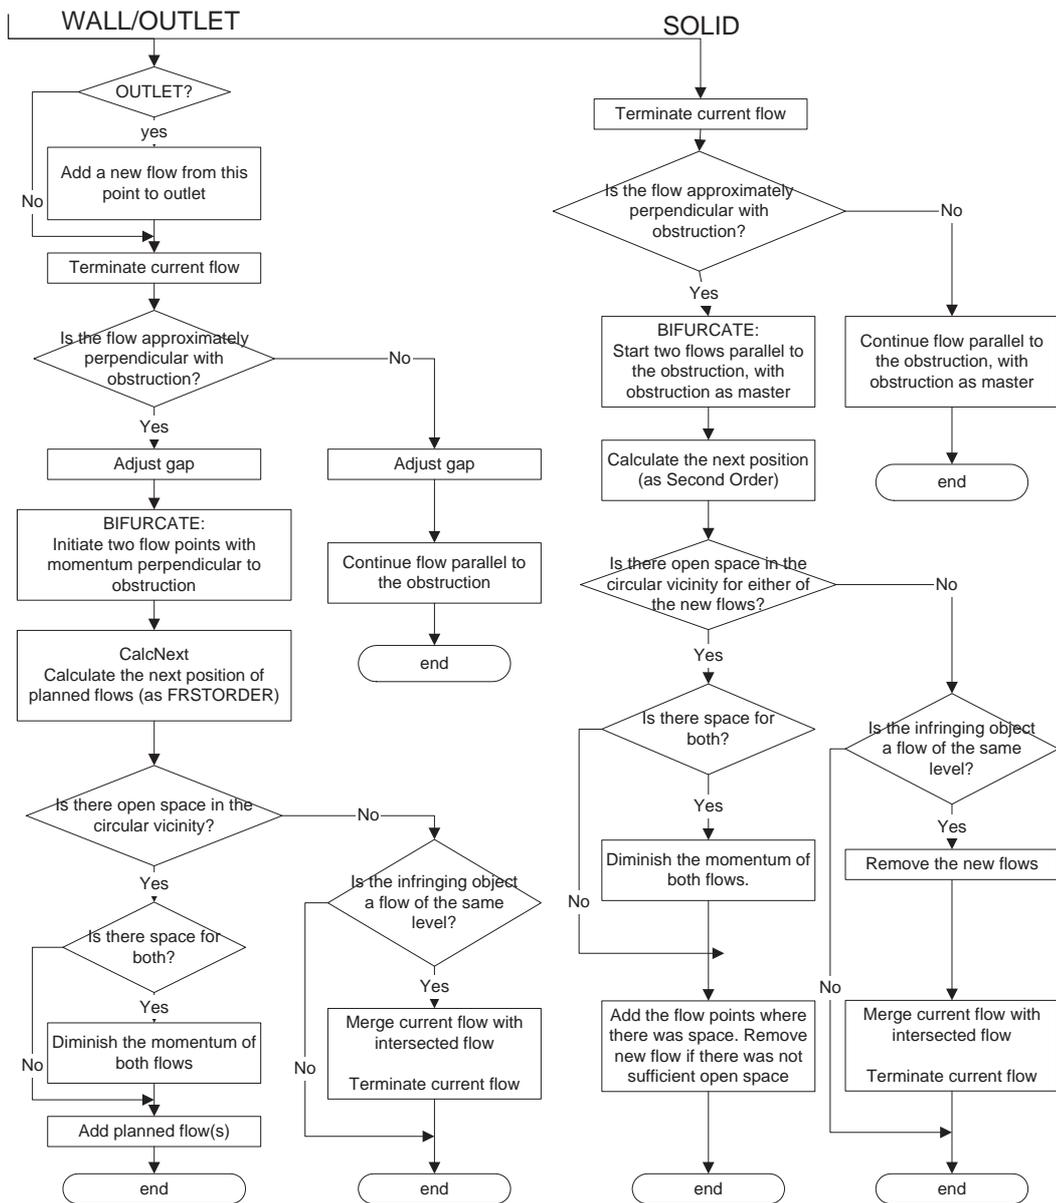


Figure 4: Increment flow: part 2

sectors begin at the start, or in fact at a flowpoint, of one of the flows that makes up the sector. The lineages are established and the sectors are detected using the parent.setAsLeaf(child) algorithm (simplified):

- If there are no lineages, add base flow and its first flow point as a lineage.
- For each lineage:
 - For each flow in lineage:
 - * Copy lineage as newLineage, and add flow and coordinate.
 - * If child exists in lineage, then a sector has been detected.
 - Record sector: all points from this collision point, through all subsequent points in all subsequent flows of lineage.
 - Add newLineage to list of lineages

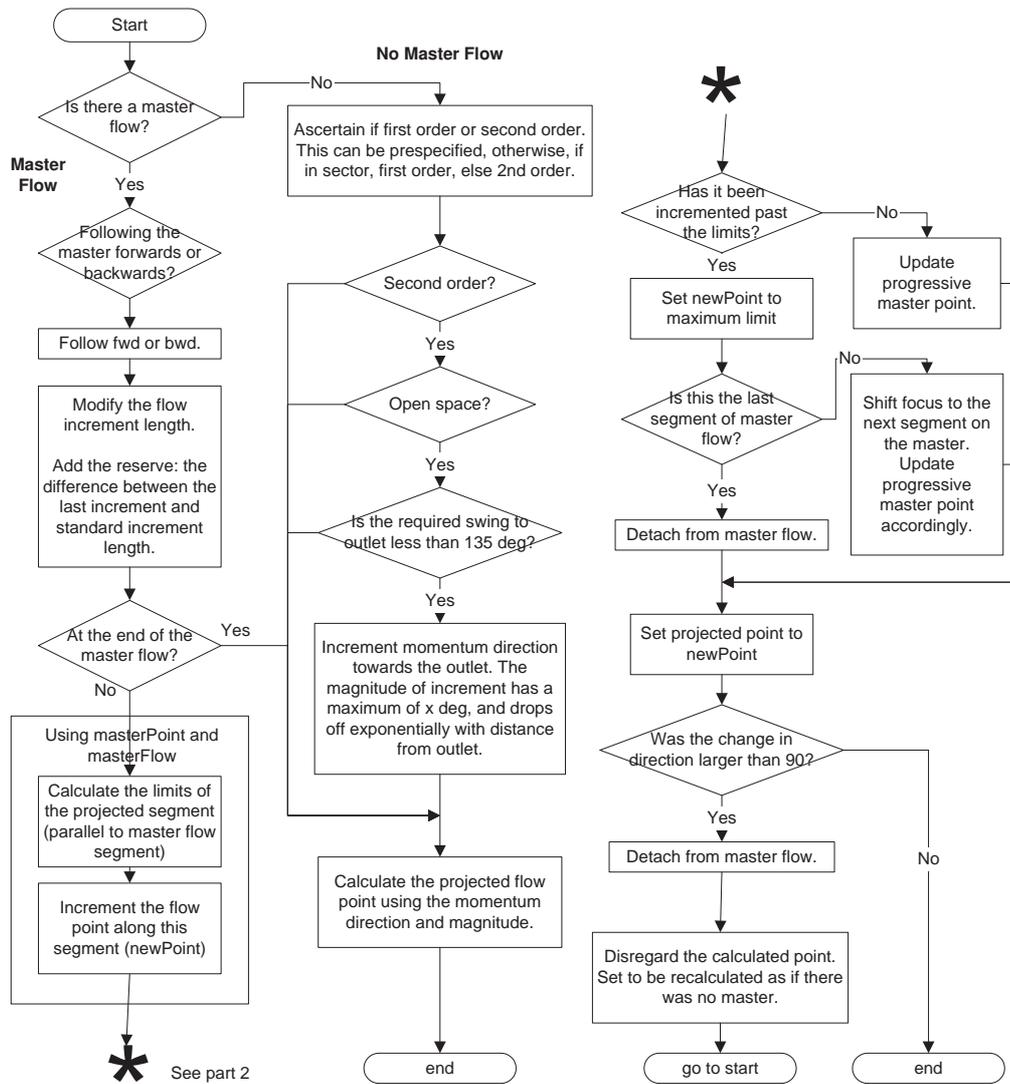


Figure 6: Calculate Next Point

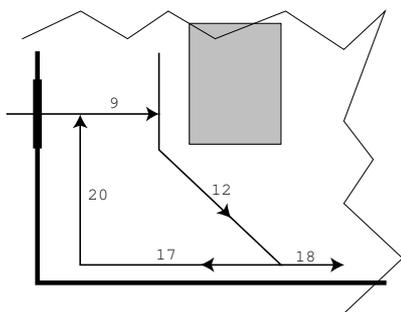


Figure 7: Flow Pattern

So that each time a new child is added, it gets a new lineage, for each existing lineage, with the same members, plus this child. The result is shown in Figure 8.

4 Results

The NaReM has been successful on a variety of configurations of inlet/outlet and object positions, number of boxes, and shape of the surroundings. A sample of result are shown in Figures 9, 10 and 11. The Flo++ fluid dynamics simulations are shown in (a) and the NaReM output shown in (b). The NaReM output displayed is the high level sector information. An example of the complete airflow model is shown in [Kowadlo and Russell, 2004].

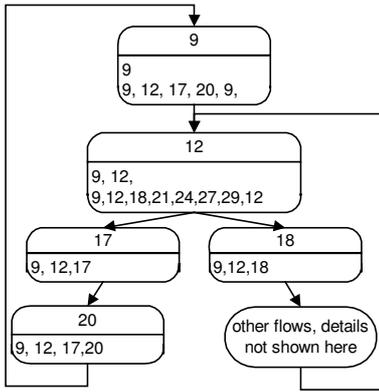
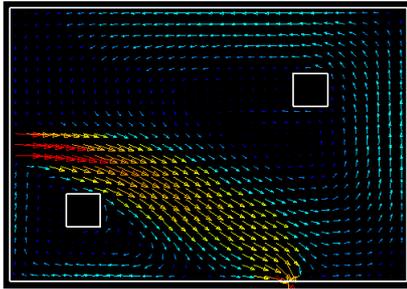
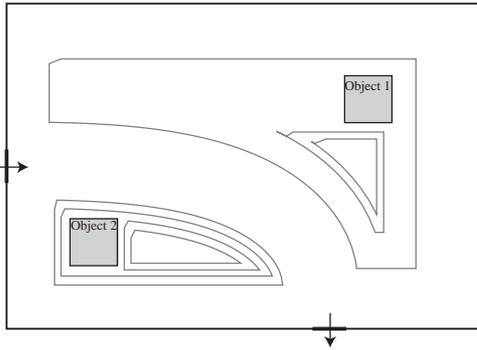


Figure 8: Flow Graph



(a) flo++



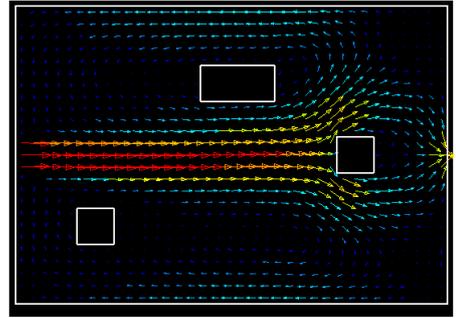
(b) NaReM

Figure 9: Scenario 1

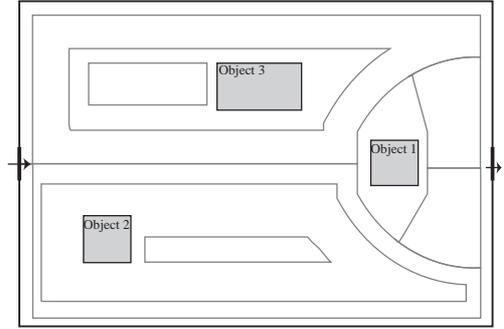
Figure 9 shows a flow being attracted to an outlet on a perpendicular wall. Figure 10 illustrates an airflow being affected by an object, changing direction but curving around it, and Figure 11 shows the way that a flow can be affected by an object by differing degrees, depending on how long it has been flowing along its surface.

5 Discussion and Future Work

The main limitation of the NaReM can be seen in Figure 12, which shows the actual flow patterns with differing vertical positions of two types of obstacles. The macroscopic flow pattern changes significantly, with even small

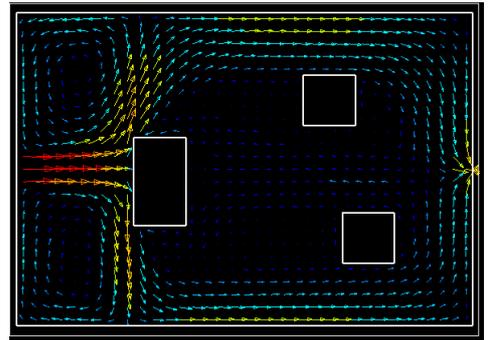


(a) flo++

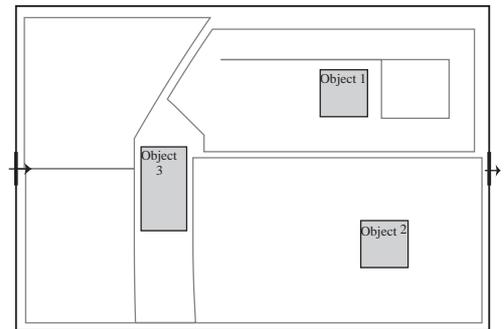


(b) NaReM

Figure 10: Scenario 2



(a) flo++



(b) NaReM

Figure 11: Scenario 3

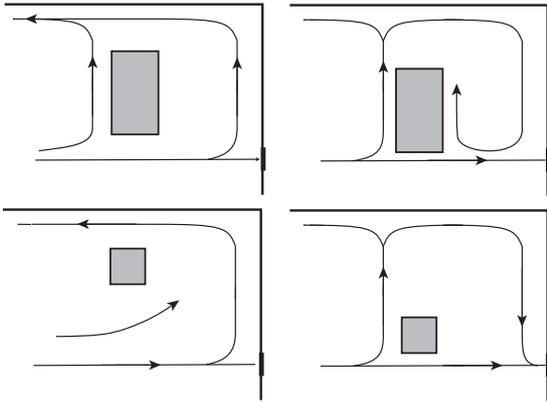


Figure 12: Flow Patterns

changes in the vertical position of the object obstruction. This is due to the flow spreading out as it propagates.

The NaReM does not predict this. The flow is propagated until completion before the rule of spreading out is implemented. Therefore, the patterns shown in the first column will be the result even if the objects are placed in the positions depicted in the second column.

An improved NaReM is currently being developed to overcome this issue. It works by generating multiple airflow map hypotheses. The area surrounding the leading edge of the flow is examined. If there is an obstacle in the general vicinity, then a new hypothesis is initiated and added as a child to the prior hypothesis, into a tree data structure, with the flow in the new hypothesis affected by this obstacle. The end result is a tree of hypotheses. Probabilities can be attached to the hypotheses depending on the distance to the obstacles, as well as uncertainty in the position of the obstacles. A hypothesis can be chosen with a combination of probabilities and information gathering (measuring airflow and chemical concentrations at key locations). This scheme has been chosen above concurrent flow spreading, as it is more robust, being able to accommodate uncertainty in the physical map, and the large imprecision of the naive model.

6 Conclusions

Using a methodology developed in past research, a NaReM or Naive Physics Reasoning Machine has been created to construct an airflow model in an enclosed environment. It has been successful at capturing the salient features of airflow for different configurations of an enclosed environment with up to three differently shaped objects, an inlet and an outlet. The model has been used as an integral part of a robot odour localisation system.

Acknowledgments

The work described in this paper was supported by the Australian Research Council funded Centre for Perceptive and Intelligent Machines in Complex Environments as well as grants from the Monash University Harold Armstrong Memorial Fund and Monash Engineering Faculty Grant Scheme.

References

- [Boxer, 2001] Paul A. Boxer. Towards learning naive physics by visual observation: Qualitative spatial representations. 2001.
- [Davis, 1997] Ernest Davis. The naive physics perplex. *AI Magazine*, 19(4):51–79, 1997.
- [Dohnal, 1991] Mirko Dohnal. A methodology for common-sense model development. *Computers in Industry*, 16(2):141–158, 1991.
- [Egenhofer and Mark, 1995] Max J. Egenhofer and David M. Mark. Naive geography. pages 1–15, 1995.
- [Forbus, 1980] Kenneth D. Forbus. Spatial and qualitative aspects of reasoning about motion. pages 170–173, 1980.
- [Gardin and Meltzer, 1989] Francesco Gardin and Bernard Meltzer. Analogical representations of naive physics. *Artificial Intelligence*, 38(2):139–159, 1989.
- [Genovese *et al.*, 1992] V. Genovese, P. Dario, R. Magni, and L. Odetti. Self organizing behavior and swarm intelligence in a pack of mobile miniature robots in search of pollutants. In *Proceedings of the 1992 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS '92)*, volume 3, pages 1575–1582, 1992.
- [Grasso *et al.*, 2000] Frank W. Grasso, Thomas R. Consi, David C. Mountain, and Jelle 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, 1979] J. Hayes, Patrick. The naive physics manifesto. In D. Michie, editor, *Expert Systems in the Microelectronic Age*. Edinburgh University Press, 1979.
- [Ishida *et al.*, 1996] H. Ishida, T. Nakamoto, and T. Moriizumi. Study of odor compass. In *Multisensor Fusion and Integration for Intelligent Systems, 1996. IEEE/SICE/RSJ International Conference on*, pages 222–226, 1996.
- [Ishida *et al.*, 2004] H. Ishida, H. Tanaka, H. Taniguchi, and T. Moriizumi. Mobile robot navigation using vision and olfaction to search for a gas/odor source. In *Intelligent Robots and Systems, 2004. (IROS 2004)*.

- Proceedings. 2004 IEEE/RSJ International Conference on*, volume 1, pages 313–318 vol.1, 2004.
- [Kowadlo and Russell, 2003] G. Kowadlo and R.A. Russell. Nave physics for effective odour localisation. In *Australian Conference on Robotics and Automation*, Brisbane, 2003.
- [Kowadlo and Russell, 2004] G. Kowadlo and R.A. Russell. To naively smell as no robot has smelt before. In *Robotics, Automation and Mechatronics, 2004 IEEE Conference on*, volume 2, pages 898–903 vol.2, 2004.
- [Kuipers, 1994] Benjamin Kuipers. *Qualitative Reasoning: Modeling and Simulation with Incomplete Knowledge*. The MIT Press, Cambridge, London, 1994.
- [Lifschitz, 1998] Vladimir Lifschitz. Cracking an egg: An exercise in commonsense reasoning. In *Common Sense 1998*, January 1998.
- [Loutfi *et al.*, 2004] A. Loutfi, S. Coradeschi, L. Karlsson, and M. Broxvall. Putting olfaction into action: using an electronic nose on a multi-sensing mobile robot. In *Intelligent Robots and Systems, 2004. (IROS 2004). Proceedings. 2004 IEEE/RSJ International Conference on*, volume 1, pages 337–342 vol.1, 2004.
- [Marques *et al.*, 2003] L. Marques, N. Almeida, and A.T. de Almeida. Olfactory sensory system for odour-plume tracking and localization. In *Proceedings of IEEE Sensors 2003, 22-24 Oct. 2003*, volume 1, pages 418–23, Toronto, Ont., Canada, 2003. IEEE.
- [Morgenstern, 2001] Leora Morgenstern. Mid-sized axiomatizations of commonsense problems: A case study in egg cracking. *Studia Logica*, 67(3):333–384, 2001.
- [Rozas *et al.*, 1991] R. Rozas, J. Morales, and D. Vega. Artificial smell detection for robotic navigation. In *Fifth International Conference on Advanced Robotics, Robots in Unstructured Environments (ICAR '91)*, volume 2, pages 1730–1733, 1991.
- [Russell, 1998] R.A. Russell. An odour sensing robot draws inspiration from the insect world. In *Bioelectromagnetism, 1998. Proceedings of the 2nd International Conference on*, pages 49–50, 1998.
- [Russell, 2001] R.Andrew Russell. Tracking chemical plumes in constrained environments. *Robotica*, 19(4):451–458, 2001.
- [Russell, 2004] R.A. Russell. Locating underground chemical sources by tracking chemical gradients in 3 dimensions. In *Intelligent Robots and Systems, 2004. (IROS 2004). Proceedings. 2004 IEEE/RSJ International Conference on*, volume 1, pages 325–330 vol.1, 2004.
- [Shanahan, 2004] Murray Shanahan. An attempt to formalise a non-trivial benchmark problem in common sense reasoning. *Artif. Intell.*, 153(1-2):141–165, 2004.
- [Smith and Casati, 1994] Barry Smith and Roberto Casati. Naive physics: An essay in ontology(1). *Philosophical Psychology*, 7(2):225–244, 1994.
- [Watt, 1995] S. Watt. A brief naive psychology manifesto. 1995.