Putting Olfaction into Action: Using an Electronic Nose on a Multi-Sensing Mobile Robot

Amy Loutfi, Silvia Coradeschi, Lars Karlsson, Mathias Broxvall
Örebro University
Örebro, SWEDEN
Web: www.aass.oru.se

***Abstract*— Olfaction is a challenging new sensing modality for intelligent systems. With the emergence of electronic noses it is now possible to detect and recognise a range of different odours for a variety of applications. An existing application is to use electronic olfaction on mobile robots for the purpose of odour based navigation. In this work, we introduce a new application where electronic olfaction is used in cooperation with other types of sensors on a mobile robot in order to acquire the odour property of objects. The mobility of the robot facilitates the execution of specific perceptual actions, such as moving closer to objects to acquire odour properties. Additional sensing modalities provides the spatial detection of objects and electronic olfaction then acquires the odour property which can be used for discrimination and recognition of the object being considered. We examine the problem of deciding when, how and where the e-nose should be activated by planning for active perception. We investigate the use of symbolic reasoning techniques in this context and consider the problem of integrating the information provided by the e-nose with both prior information and information from other sensors (e.g., vision). Finally, experiments are performed on a mobile robot equipped with an e-nose together with a variety of sensors that can perform decision making tasks in realistic environments.**

I. Introduction

Mobile robots are becoming equipped with more numerous and more diverse sensing modalities. One recent modality that is attracting increasing interest is the ability to integrate electronic olfaction on mobile systems. The development of compact and commercial gas sensors has facilitated this integration [3]. To date, the research areas that have used electronic olfaction on mobile robotic systems have concentrated mainly in the applications of odour source location or odour source tracking. That is to say, to navigate a robot according to its perceptions of smell . Although these works provide an interesting study on different search strategies [5], they have yet to utilize the full capabilities of electronic olfaction. This is mainly due to the fact that experiments are conducted with only one type of odour source and use gas sensors which are specifically tuned for the odour in question. Furthermore, from a conceptual point of view, the use of smell alone cannot detect and localize physical objects. Conceptually, smell can be used only to describe odour character. Thus, in order to do proper object detection, other sensors are required that can provide either spatial features or tactile features.

The novelty of this work is to combine vision and smell on a mobile robot in order to detect objects and associate the smell property to the object. In this work, we also consider a variety of odours and take advantage of the e-nose's ability to distinguish between known odours as well as acquire and learn information about new odours on-line. Our ambition is to integrate the e-nose in a mobile platform that is able to combine the processes of decision-making, acting and sensing. A naïve approach may be to use an electronic nose as a passive sensor i.e., constantly smelling the environment and associating the odour characteristic to the object in focus. There are however, several problems to passive sensing in this case. First, there is the conceptual problem that questions the validity of associating a dispersed odour in the air to an object physically located at a distance from the actual point of detection. There is also the practical problems of the sensing mechanism that include long sampling time (1-3 minutes), high power consumption (pumps and heaters), and long processing time for the multivariate sensor data. In a real time application that considers a mobile platform with multiple sensing modalities, the electronic nose is an expensive sensor. For these reasons, the electronic nose in this work act as an active sensor that is explicitly called upon inside a complex decision making system. In order to effectively execute the active sensing, a planner is integrated into the system to plan for specific perceptual actions. Although planning for perceptual actions using active perception has been previously considered, experiments have only focused on either single sensing modalities or vision based modalities [1] [4] [9]. Similarly, sensor fusion techniques [10][7], have yet to consider the combination of different sensors with an electronic noses. By integrating different sensors of various types, the planner needs to decide which sensors should be used, when they should be used, and how to act in order to resolve ambiguous cases.

The contribution of this work can be summarized in two points. First, we consider the use of electronic olfaction on a comprehensive system that includes decision making, symbolic reasoning and different sensing modalities on a mobile platform. It is believed that such a system can benefit a variety of applications ranging from the detection of hazardous objects in airport security to fast localization of odour leaks related to physical objects (leaking gas

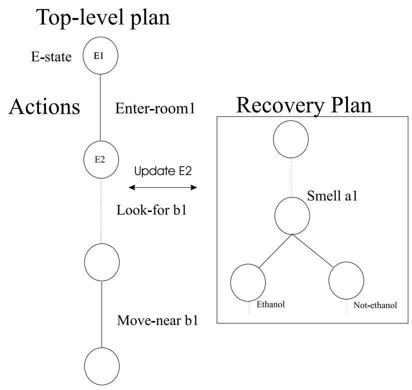

Fig. 2. Relation between *top-level* and recovery planning

bottles), and even in the discrimination of visually identical objects. The secondary contribution lies in the more general framework of multiple-sensing integration on embedded systems. We show that autonomous systems need to not only manage the information from multiple modalities but also need to effectively employ these modalities in order to acquire reliable perceptual information. We show that explicit planning for perceptual actions can be used in the presence of sensors that either acquire different perceptual properties and/or expend various resources.

The paper begins with an overview of the system and its components in Section 2. In Section 3, the actual integration of the electronic nose with the system and how active perception is implemented in a dynamic situation is explained. Finally, in Section 4, different scenarios are tested using vision together with olfaction on a robotic platform.

II. System Architecture

The components of the complete system used in our experiments and their interrelations are shown in Figure 1. Significant attention is placed on the olfactory module with a description of the sensor operation and data processing.

A. Planner

The planner, PTLplan, is a conditional progressive planner [6]. It has the capacity to:

- reason about incomplete and uncertain information
- reason about informative actions
- generate plans that contain conditional branches

The planner functions by searching in a space of epistemic states, or e-states. An e-state represents the agent's incomplete and uncertain knowledge about the world at some point in time. An e-state can be considered to represent a set of hypotheses about the actual state of the world. The planner can reason about perceptive actions and these actions have the effect that the agents makes observations that may help it to distinguish between different hypotheses. Each different observation will result in a separate new and typically smaller e-state, and in each e-state the robot will know more than before.

Figure 2 shows how the PTLplan is used on two different levels: for constructing the plan for achieving the current top level goal, and for constructing recovery plans in case problems occur while the main plan is executed. The planning process initially begins with the execution of a top level plan. If and when a perceptive action in the plan involving the anchoring module fails, a recovery planner is invoked. Failures may arise either because several candidates partially matching the description of the requested object are found (ambiguity), or no candidates are found. The recovery planner functions by collecting information about the requested object from the current e-state in the top level plan and adds information from the anchoring module about actual observation from the candidate object. Based on this information, the planner creates a recovery e-state with a number of different hypothesis of which objects are the actual candidate objects. A recovery plan is generated with the goal to determine which of the hypotheses is the correct one. The actions in the recovery plan are chosen from a restricted action set (e.g. move near objects, smell objects) that do not alter the current situation too much. In the figure an example is illustrated where the top-level-plan calls the following actions: room1 - look-for b1 - move-near b1. In this case the action to look-for b1 fails (because of an ambiguity between b1 and a1) and the recovery planner is activated which generates (touch a1, smell a1, cond ((wine) :success) ((not-wine) touch a2, smell a2, cond ...)). The new actions of the recovery plan are executed and the top level e-state is updated.

B. Plan executor

The plan executor takes the individual actions of the plan and translates them into tasks that can be executed by the control system (the Thinking Cap). These tasks typically consist of navigation from one locality or object to another. Planning actions might also be translated into requests to the anchoring system to find or reacquire an object that is relevant to the plan, either in order to act on that object (e.g. moving towards it) or simply to acquire more information about the object. The plan executor also reacts when the execution of an action fails, e.g. due to ambiguities when it tries to anchor an object. In such cases, the recovery planning facility is invoked. If the recovery plan is completed sucessfully, execution of the failed top-level action is resumed.

C. Anchoring Module

The anchoring module creates and maintains the connection between symbols referring to physical objects and sensor data provided by the vision and olfaction sensors. The symbol-data correspondence is represented by a data structure called an anchor, that include pointers to both the symbol and the sensor data connected to it. The anchors also maintain a set of properties that are useful to re-identify an object e.g., colour and position. These properties can also be used as input to the control routines. Different functionalities are included in the anchoring

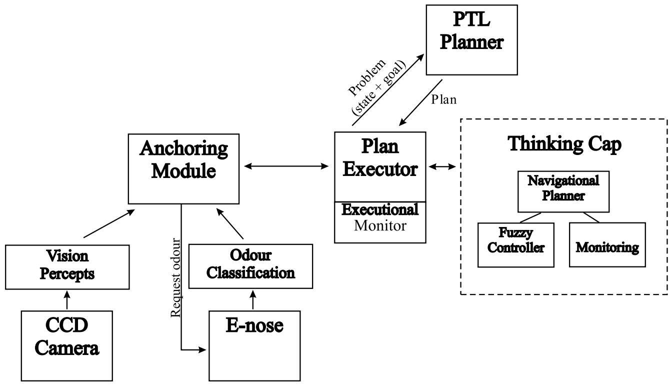

Fig. 1. Overview of the complete system that anchors symbols using olfaction and vision percepts

modules. In this work, two functionalities in particular are used.

Find - is used to link the symbol e.g., "bottle-22" to a percept such as a region in an image that matches the description "red bottle containing wine". The output of Find is an anchor that contains properties such as the *(x,y)* position or the odour of the bottle.

Reacquire - is used to update the properties of an existing anchor. This may be useful if the object goes out of view or a period of time elapses resulting in a change of object properties (e.g., chemical characteristic).

The anchoring functionalities are typically called by the planner via the plan executor. To be able to execute actions referring to an object, the planner interacts with the anchoring module by referring to objects using a symbolic name and a description expressed in terms of predicates. For instance, we can execute the command "move-near bottle-25" where"bottle-25" is described as a "green bottle".

Since all properties of an object are not always accessible, the anchoring module also considers cases of *partial matchings*. We consider a matching between a description and the perceptual properties of an object partial when all perceived properties of the object match the description, but there still remains properties in the description that have not been perceived. This is a typical case for olfaction that requires that the sensors are close to the odour source for detection.

The anchoring module is also able to determine whether objects have been previously perceived, so as to not create new anchors for existing objects. Ambiguous cases such as when two objects partially match a given description, and failure to find an object are detected by the module and dealt with at the planner level.

D. Thinking Cap

In this system, execution monitoring on a mobile robot is controlled by a hybrid architecture evolved from [8] called the Thinking Cap. The Thinking Cap (TC) consists of a fuzzy behaviour-based controller, and a navigation planner. In order to achieve a goal the planner selects a number of behaviours to be executed in different situations. Depending on the current situation, the different behaviours are activated to different degrees. The Thinking Cap behaviours are continuously active, this means that even when actions are being performed from the planner, certain routines such as obstacle avoidance, and keeping off of objects are maintained.

E. Vision Module

In addition to the sonars used by Thinking Cap to navigate and detect obstacles, the system also uses vision to perceive objects. This is done by continuously receiving images from a CCD camera connected to the robot and using standard image recognition techniques for image segmentation. The segmented images are used for recognising a number of predetermined classes of objects and properties such as shape, colour and relative position. The resulting classified objects are delivered to the rest of the system at approximately 1fps.

The result of these classifications are collected over time and presented to the planning and anchoring system. The system tries to represent objects so that they are persistent over time but due to uncertainties in our sensing this is not always possible and ambiguities which has to be dealt with at the planning level may arise.

In order to maintain and predict sensed object currently outside the camera's viewpoint an odometry based localisation is used. All perceived objects are stored in a list of trajectories which can be accessed by the anchoring module to establish anchors. For limited movements, the system can easily reacquire objects based on their stored position. If movements are large, or objects move or if accumulation of odometry errors is to big this might lead to reacquisition ambiguities which can only be resolved using other sensors. All perceived objects are stored in a list of trajectories which can be accessed by the anchoring module to establish anchors.

F. Olfactory Module

The olfactory module consists of a commercially available electronic nose. This e-nose contains 32 thin-film carbon-black conducting polymer sensors of variable selectivity. Each sensor consists of two electrical leads on

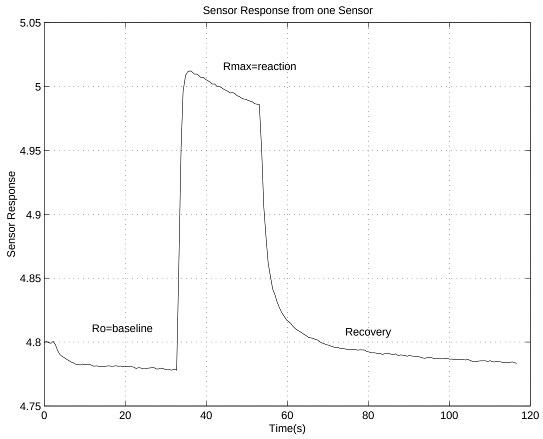

Fig. 3. Example of sensor data throughout the three phases of sampling an odour with one of the 32 gas sensors.

an aluminum substrate. Thin films are deposited across the leads creating a chemiresistor. Upon exposure to an analyte, a polymer matrix absorbs the analyte and increases in volume. This increase in volume is reflected in a increase in resistance across the sensor. Each polymer in the array is unique and designed to absorb gases to different degrees, creating a pattern of responses across the array. The array of sensors are contained in a portable unit also consisting of pumps, valves and filters that are required to expose the sensor array to a vapour gas.

The sampling of an odour occurs in three different phases. The first phase is a baseline purge, where the sensors are exposed to a steady state condition, for example the air in the room. The duration of the purge is 30 seconds. The second phase is a sampling cycle where a valve is switched to allow vapours from a sampling inlet to come into contact with the sensing array. The duration of this cycle is 20 seconds. Finally, a sequence of purging cycles is used to remove the sampled odour from the unit and restore the sensor response to the baseline values. Figure 3 illustrates a typical response from one sensor.

The signals are gathered in a response vector where each sensor's reaction is represented by

$$\frac{\Delta R}{R_o} = \frac{(R_{max} - R_o)}{R_o} \tag{1}$$

The response vectors are then normalised using a simple weighting method and autoscaled. Classification of new odours is performed by first collecting a series of training data. All odour samples are maintained in a repository that is updated on-line. For the purpose of the experiments described below, a simpler method was chosen in order to optimise computation time in a real-time environment and a minimum distance classifier was implemented. The result from this classifier provides a linguistic name which refers to the class to which the unknown odour belongs.

III. Perceptual Actions using an E-nose

To integrate the olfactory component we need to satisfy three criteria:

Active Perception - Requires the ability to perform specific actions in order to gain perceptual information. Unlike the vision modality that is continuously sampling the environment, the electronic nose is explicitly called upon when needed. This call originates from the planner that decides when the e-nose should be used and prompts the anchoring module to send the request to the e-nose. The sampling of the e-nose requires that the planner creates a multi-step plan. The first step is to approach the object by activating a “move-near” behaviour. The TC uses the position property provided by the anchor to control the robot. When the robot is sufficiently near the object the next step of the plan is executed and the e-nose is activated. Upon completion of odour sampling, the classification is sent back to the anchoring module and the anchoring module signals to the planner that the smell property has been acquired. Based on the value of the property, the planner decides how to continue the execution of the plan. Note that all processes are sent through the anchoring module since the anchors are responsible for maintaining the information regarding the symbolic properties of the objects. The anchor is also responsible for keeping track of which object is currently smelled and attributing the correct smell to the correct object. The planner needs to carefully decide when the olfactory sensor needs to be called. This is because of the eventual cost of running the nose is high in terms of time and power consumption. In the experiments performed, the planner makes the call to smell only when the goal cannot be achieved by using the spatial properties alone.

Domain Knowledge - Guides the perceptions using external information. In this case, this external information provides information regarding the names of trained odours, the thresholds for object segmentation, the goals of the system. For the olfactory module, the odour names are often correlated to the name of the substance which emits the odour e.g, ethanol, linalool. Depending on the application, the name of the odour may also reflect a concentration level or the odour character in terms of quality.

Online Learning - The results of all perceptual actions are stored and contribute to the development of a database of odours that can then be used for further recognition. A repository maintains olfactory signals from the e-nose. The collection of odours is used in training process when unseen samples are to be recognized.

IV. Experiments

A. Setup

The experimental setup consists of a Magellan Pro Research robot, called Pippi, equipped with a CCD camera, sonars, infrared sensors,a compass and a Cyranose 320 electronic nose. A snout protrudes from the robot so that sampling of an object can be done easily see Figure 4. The nose has been previously trained on a number of different substances including those used in the first experiment.

B. Disambiguating Objects

The aim of the first experiment is to show how the e-nose can discriminate objects using their odour properties. Pippi is located in a room where several objects are present. The robot can wander throughout the room using the spatial properties to reccognize known objects and to avoid obstacles. Using the list of trajectories, Pippi is able to keep track of the objects that she has previously perceived even if they go out of view of the camera. She can also re-establish their identity once they come into view again.

When Pippi is asked to find objects, she does so using the visual properties. However, cases may arise where the visual properties are not adequate to distinguish between objects. In these cases the anchoring module detects an ambiguity and informs the planner of the different candidate objects. The planner copes with this situation by invoking the recovery planner which determines which additional properties of the objects can be acquired and also determines how to acquire these properties. Calls to other sensing modalities may be made which may result in the execution of actions. Depending on the perceptual information received, the goal is either attained or the planner resumes the conditional plan.

The following experiments are performed by using a number of cups that contain different substances. The cups are scattered throughout the room and the goal of the Pippi is to find a cup that is characterized by both a visual and a odour description. Also in the room are a number of obstacles and dynamic variables such as people and other robots moving in the environment. The planner activates the find functionality creating anchors for the matching objects. The anchoring module examines the percepts sent from the vision module, and finds that several percepts match. Figure 5 shows an example when there are 4 identical cups located in the room, each object is denoted by its shape and a number that refers to its anchor value. In these kinds of situation an ambiguity is detected and a plan is generated to smell each of the cups. Table I summarizes the results from different scenarios considering different number of ambiguities. Note that in each case to execute the smell action Pippi needs to move close to the objects. As a result, errors may arise from either the olfaction module (misclassification of odours) and/or the vision module (odometry problems, looses track of an object). The table also provides information regarding the source of failures in the unsuccessful cases.

For each ambiguous case a number of different scenarios are considered. In the first case, two green cups are located in a room that contain two different substances, Ethanol and Hexanal. Pippi is asked to either find the Ethanol or the Hexanal cup in a total of 11 different scenarios. The scenarios vary in that the cups are placed in different positions, objects are re-located and environment is varied. In successive cases, the number of cups and odours are increased. Finally, five identical cups containing Ethanol, Hexanal, 3-Hexanal, Octanol, Linalool. These odours are part of a ASTM atlas of odour descriptions [2] whose character is best described as alcoholic,sour,woody, oily and fragrant respectively. The analysis of the results shows that visual failures provoke olfactory failures. This is due to the fact that e-nose performs best when close to an object. Depending on the odour (rate of vapourization), the distance to accurately recognize the odour range is between 5 cm and 23 cm. Most visual failures are due to the odometry and as the number of ambiguous cases increase, the more movement Pippi needs to perform and thus an increase in odometry error. There are cases in which the e-nose mis-classifies the odour independently of visual failures. These misclassification errors slightly increase when the number of different odours increase. The source of this error can be the e-nose's inability to discriminate between classes of odours. However, in our case the error was actually due to the sensing parameters given in the sampling process. In particular, when two cups were separated by a short distance there was an inadequate recovery time between "sniffs" and this resulted in a misclassification of samples. This recovery time depends on the type of odour being sampled.

Fig. 4. A situation where the robot discriminates similar green cup objects based on their smell property

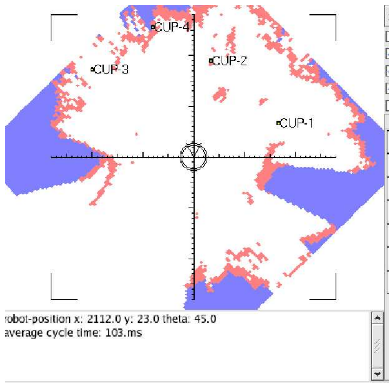

Fig. 5. The local perceptual space of Pippi given 4 identical cups and obstacles. Pippi is located in the center of the space.

C. Reacquiring Objects

In the previous experiment, the planner was able to use the e-nose to disambiguate between visually identical objects. To do this, the planner needed to obtain the

TABLE I

EXPERIMENTAL RESULTS FROM AMBIGUOUS CASES

Number of Odours	Scenarios	Visual and Olfactory Failures	Olfactory Failures	Successes
2	11	2	0	9
3	15	3	0	12
4	21	4	1	16
5	25	4	2	19

classification results from the e-nose in order to determine the content of each of the cups. In a real situation, it is difficult to predict which substances the nose will need to classify. Therefore, to be able to adapt to new odours, the e-nose needs to be able to learn online and update its odour repository. To illustrate how this is possible, a series of experiments are done using the reacquire functionality explained in the anchoring module.

In these scenarios, the robot is in a room and a cup is located on the floor. Instead of finding a specific object, Pippi is is given the task to acquire as much information about the objects in the room. She does this by perceiving each object visually, then deriving a plan to smell each object. If Pippi is able to recognize the smell, she associates that smell with the anchor. If however, recognition is unsuccessful, Pippi creates a new entry in the repository for that odour print and gives it an arbitrary name that is stored with the anchor.

Pippi is then allowed to wander in and out of the room. While she is wandering, the objects are moved and new objects are added. We request from Pippi to reacquire an object previously perceived. Because the objects have changed position and new objects are present, Pippi cannot rely only on the visual percepts to discriminate the required objects. Since several possible candidates are found, the planner is informed that there is an ambiguous situation. A plan is created consisting of first going to one of the cups and smelling it, if the classification of the odour matches the one that was stored during the first acquire then the plan succeeds. Otherwise, the other objects are checked. The results from the reacquire experiments yielded a similar pattern in terms of success as the previous experiments. Most failures were attributed to vision and odometry which propagated into olfactory failures.

V. CONCLUSIONS

Olfaction is a valued sense in humans, and robotic systems, which interact with humans, and/or execute human-like tasks also benefit from the ability to perceive and recognise odours. While previously gas sensors were difficult to use and needed certain expertise to successfully implement odour recognition, commercial products have now made it possible to successfully employ electronic olfaction in new domains. One such domain is intelligent systems that rely on multi-sensing processes to perform autonomous tasks. The integration of electronic olfaction presents interesting challenges with respect to the use of AI techniques in robotic platforms. Some of these challenges are due to the properties of the sensing mechanism, such as long sampling time, and close proximity required for smelling an object.

In this work, we show how an electronic nose could be successfully used as a tool for the recognition of objects. We also show that intelligent techniques can be implemented to manage sensing modalities and to use these modalities appropriately. Specifically we focus on the use of symbolic techniques, which have allowed the planning of perceptual tasks. Experimental results were able to illustrate how a mobile robot with multiple-sensing modalities at its disposal could cope in different situations that required both olfactory and spatial perceptions. These types of experiments are not intended to replace odour based navigational research, rather they compliment and emphasize the broad potential of using olfaction on mobile systems.

REFERENCES

[1] C. Barrouil, C. Castel, P. Fabiani, R. Mampey, P. Secchi, and C. Tessier. Perception strategy for a surveillance system. In *Proc. of ECAI*, pages 627–631, 1998.

[2] A. Dravnieks. *Atlas of Odor Character profiles (ASTM Data Series Publication DS 61)*. American Society for Testing, USA, 2000.

[3] J. Gardner and P. Bartlett. *Electronic Noses, Principles and Applications*. Oxford University Press, New York, NY, USA, 1999.

[4] S. Kovacic, A. Leonardis, and F. Pernus. Planning sequences of views for 3-D object recognition and pose determination. *Pattern Recognition*, 31:1407–1417, 1998.

[5] A. Lilienthal, A. Zell, M. Wandel, and U. Weimar. Sensing odour sources in indoor environments without a constant airflow by a mobile robot. In *Proc. of the IEEE Int. Conf. on Robotics and Automation (ICRA)*, pages 4005–4010, Seoul, South Korea, 2001.

[6] L.Karlsson. Conditional progressive planning under uncertainty. In *Proc. of the 17th Int. Joint Conferences on Artificial Intelligence (IJCAI)*, pages 431–438, 2001.

[7] R. Murphy. Allocating sensing resources with incomplete information. In *AAAI Spring Symposium*, 1996.

[8] A. Saffiotti, K. Konolige, and E. H. Ruspini. A multivalued-logic approach to integrating planningand control. *Artificial Intelligence*, 76(1-2):481–526, 1995.

[9] G. Wasson, D. Kortenkamp, and E. Huber. Integrating active perception with an autonomous robot architecture. In *Proc. of the 2nd Int. Conf. on Autonomous Agents (Agents)*, pages 325–331, 1998.

[10] H. Xu and J. Vandorpe. Perception planning in mobile robot navigation. In *Proceedings of the IEEE Internation Conference on MFI '94*, pages 723–730, 1994.