Improving Odour Analysis Through Human-Robot Cooperation

Amy Loutfi and Silvia Coradeschi
Center for Applied Autonomous Sensor Systems
University of Örebro
Örebro, Sweden
www.aass.oru.se

***Abstract*— More and more work in the field of artificial olfaction considers the integration of olfaction onto robotic systems. An important part of this integration is providing the robot with the ability to discriminate between different odour substances. This work presents the integration of an electronic nose onto a complete robotic system with multi-sensing modalities. The ambition of this work is to illustrate how classification performance of odours can be improved through the exploitation of the mobility of a robot, as well as through the cooperation between a human user and an electronic perceptual system. These points are highlighted in the context of a online robotic system designed to perform different tasks which require the ability to discriminate between odour characters. The robotic system and experimental results are presented where these tasks are tested and evaluated.**

***Index Terms*— electronic olfaction, human-computer interface, anchoring, planning for perceptual actions**

I. Introduction

The ability to mimic biological olfaction using artificial sensors has been an area of continuous interest since the emergence of the compact gas sensor in the late 20th century. To date applications that involve artificial olfaction include quality evaluation in the food industry, environmental monitoring, and detection of obnoxious and hazardous odours. Many of these applications use what is called an electronic nose. That is to say, an array of gas sensors each of varying selectivity along with a pattern recognition component trained to discriminate between both simple and complex odours [3].

Today, a handful of robotic systems include artificial olfaction. However, these systems mainly include simple gas sensors used for odour-based navigation often where a priori knowledge of the target odour is available[1] [6], [11], [13]. There still remains the challenge to integrate odour analysis and classification on robotic, mobile and autonomous platforms. Eventually, as robots become more integrated in our environment and equipped with greater diversity of sensors, the ability to perform the analyse of an odour characteristic would add considerable advantages to many applications. Additionally, the ability to communicate the results to a human user could be equally appreciated. Potential applications that could benefit from these two components include, robots in cleaning, rescue, space exploration, health care and mining.

In our previous work [10], we presented a system where a multi-sensing robot integrated odour discrimination to assist in object recognition tasks. We showed that this integration was in fact a non-trivial task whereupon several sensing modules are required to cooperate in order to achieve a specific goal. We illustrated that the electronic nose functions best as an active sensor, explicitly called upon by an autonomous system capable of making decisions, navigating by vision and assessing uncertainties. In this paper, we aim to extend the previous work to illustrate that not only can a mobile robot system benefit by using odour analysis, but that the mobility and autonomy of the robot can also be exploited to better the odour analysis. A second and equally important novelty of this work is contained in a human-computer interface specifically tailored for a robotic system capable of discriminating between different odour characters.

A. Our Approach

Our approach is motivated by a larger vision to have an online robotic system working for longer periods of time in a dynamic environment. The system should be able to use its ability to discriminate odours to perform different odour related tasks, as well as improve its knowledge about odours as the state of the world changes. Certain criteria were considered in the design of this system, this criteria was based on the properties of the olfactory sensing module, the intention of the system and the interaction with human users. Consequently, the system architecture attempts to satisfy these criteria by considering the following:

1) A robot with its own representation of odours: Electronic noses can be sensitive to some odours undetectable in the human perceptual domain, whilst insensitive to other very common odours. Most electronic noses consists of large array of sensors whose sampling frequency is high and exposure time long resulting in a difficult task for multi-data analysis. Thus the data is often unintuitive. To be able to obtain an accurate representation of the odors

[1] Few works have also considered building concentration maps [7] but still the discriminatory analysis between different odour components have yet to be explored.

that corresponds to the actual sensitivity of the e-nose, the electronic nose needs to be able to create its own representation using unsupervised algorithms. This has the advantage to let the e-nose use its full potential and to give the possibility to the user to recognize weakness and strengths of the e-nose. This knowledge can be used to better evaluate the results of the e-nose, but also to tune the e-nose toward specific application by selecting appropriate sensors.

2) A symbolic representation of perceptions and tasks: The representation of information on a symbolic level is advocated for several reasons. First, most tasks involving odour recognition presented in this work requires some degree of object recognition. For example, the ability to identify a coffee cup, or a suspicious package. To accomplish this task, different sensor modalities are used. The use of symbols allows us to summarize sensory information relating to the same physical object into a more structured representation called anchors, even when this information is available at different times. Secondly, odour classification in general relies heavily on symbolic, specifically linguistic terms. As presented in [2], odour classification lacks a standard metric most likely attributed to the absence of psycho-schematic representation (the names of many common odours have little correspondence to their chemical composition). Consequently, in the human perceptual domain language is instrumental in communicating perceptions. A system whose ambitions are to interact with humans (experts and non-experts) needs to be able to convey perceptual information in a manner that can be understood.

3) A human-machine interface: The robot that we envision performs its tasks autonomously, but under the direction of a human selecting the tasks and helping improving the odor classification. An interesting challenge to the odor classification is the possibility of conflict between the two objectives outlined above. Since maintaining the sensory interpretations of the e-nose and integrating symbolic representation that is understood by a human may cause difficulties. An interface where upon a human can interact with the robot can help in resolving the conflict. Symbols can be edited to give more meaning in the human perceptual domain and the user can clearly evaluate the capability of the robot.

The paper is organized as follows: Section II provides an overview of the robotic system architecture. Section III discusses the human-computer interface used to control the robot. Section IV presents in detail key issues to consider when performing odour analysis, and presents experimental results that show how a human and robot can cooperate to solve odour related tasks. Finally, conclusions and future works are presented in Section V.

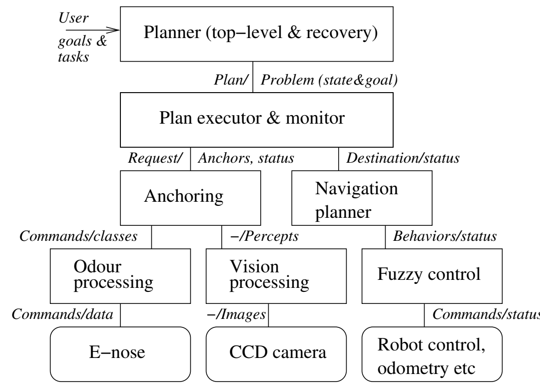

Fig. 1. Overview of the complete system. On the left side of the arrows information is flowing downward and on the right side information is flowing upward.

II. System Architecture

A. Olfaction

The olfactory module shown in the complete system (Figure 1), consists of a commercially available electronic nose. This e-nose contains 32 thin-film carbon-black conducting polymer sensors of variable selectivity [5]. Each sensor consists of two electrical leads on 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. 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 2 (*Left*) illustrates a typical response from one sensor.

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

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

and a symbolic tag, L_o, provided by the anchoring module (Section II-D). The response vector for an odour, also called a smell print is shown in Figure 2 (*Right*).

In this work we implement a semi-supervised fuzzy-clustering algorithm. The use of a fuzzy algorithm provides for the explicit representation of uncertainty in the classification output (this uncertainty can later be used by the planner to initiate further perceptual actions); it also provides for easy integration of symbolic representation without the need to map specific outputs to categorical names. Another benefit to fuzzy clustering is its suitability for small sample sizes. The algorithm applied is based on a combination of the fuzzy c-means and a fuzzy maximum likelihood estimator (FMLE) [4]. The number of clusters, n, is given by the human user as the number of expected odours, however, no information is given with regards to which samples belong to which cluster. Clusters centers are determined by the properties of the sensor data alone, using each odour's response vector as input. To map each cluster to an odour character, a contingency table represents the probability of an odour label L_o occurring within a prototypical region around the cluster center. Prototypical regions are defined as a region on the data space where membership to a cluster is equal to 1. The membership degree is calculated using:

$$u_{ik} = \frac{1}{\sum_{j=1}^{c} \left(\frac{d^2_{F_j}(x_k, v_i)}{d^2_{F_j}(x_k, v_j)} \right)^{\frac{2}{m-1}}} \qquad (2)$$

where $d^2_{F_j}(x_k, v_i)$ is the FMLE distance between the k^{th} feature vector and the i^{th} cluster center and m is a fuzzification factor set to 2.

Labels which most frequently occur within a particular cluster center are then used as the naming convention for that cluster. It is in this way, that the electronic perception of odours is preserved meanwhile, coordinating the symbolic representation of odours used by a human. More detail on the odour classification can be found in the authors' works [9].

Two methods are used to evaluate the classification performance, uncertainty coefficient and validity measures. The coefficient of uncertainty measures the entropy in the contingency table. Values closer to 1 represent a strong correlation between the human perception of odours and the electronic one, whereas value closer to 0 indicate little correlation. The validity measures, evaluate the properties of the sensor data by qualifying the clusters from a global aspect. In this work we value compact, dense and well separated clusters. We use a Compactness and Separation [14] which minimizes for compact and well separated clusters.

B. Vision

In addition to the sonars used to navigate and detect obstacles, the system also uses vision to perceive objects.

In order to maintain and predict sensed objects currently outside the camera's viewpoint an odometry based localization is used so that objects not currently perceived can be remembered. 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, objects move or if the accumulation of odometry errors is too large this might lead to reacquisition ambiguities which, in our case are resolved by engaging a recovery planner (Section II-E).

C. Plan Execution and Monitoring

Execution monitoring on a mobile robot is controlled by a hybrid architecture evolved from [12] 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.

D. Anchoring

The anchoring module functions to create and maintain the connection between symbols and symbolic descriptions referring to physical objects and the objects perceived by the sensors. The symbol-data correspondence is represented by a data structure called an anchor [1], that includes a pointer to both the symbol and the sensor data connected to it. 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.

E. Planner

PTLplan is a planner for partially observable domains with uncertainty (probabilities) [8]. It searches in a space of epistemic states, or e-states for short, where 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. Each sub-state in the e-state consists of a set of value-assignments to properties or state variables, where a property might represent for instance the shape or smell of a particular object: (shape b1 = cup), (smell b1 = ethanol). Sub-states are also associated with degrees of probability. The planner can reason about perceptive actions, such as looking at or smelling an object, and these actions have the effect that the agent makes observations that may help it to distinguish between the different hypotheses. Each different observation will result in a separate new and typically smaller e-state, and in each such e-state the agent will know more than before. Figure 3 illustrates the effects on e-states of two actions: one for movement, and one for perception (smell). It would also be possible to represent sensor errors — observations need not always coincide with the actual state.

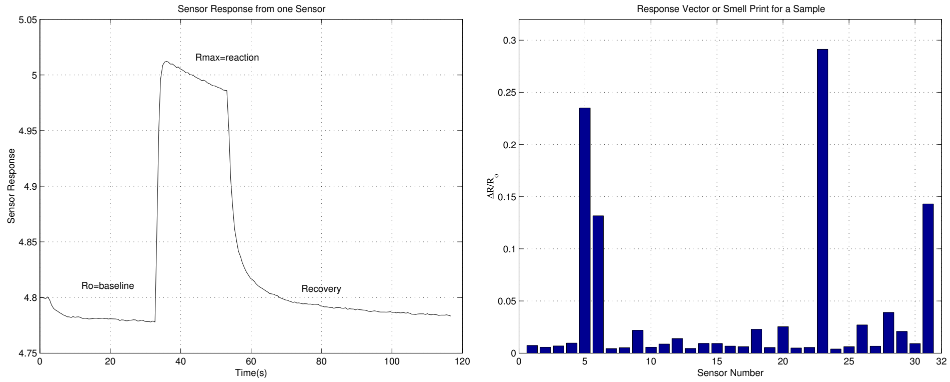

Fig. 2. *(Left)* Example of sensor data throughout the three phases of sampling an odour with one of the 32 gas sensors. *(Right)* The smell print for an odour showing the response vectors over the entire sensing array from Equation 1

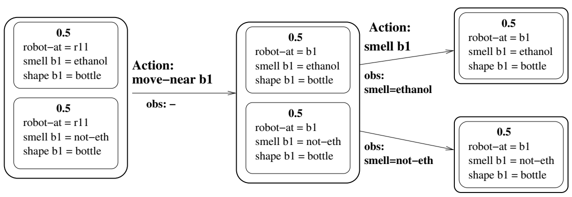

Fig. 3. Epistemic states (thick borders) contain one or more sub-states with associated probabilities (hypotheses), and are transformed by the applications of actions (arrows). Note how different observations yield different new e-states.

PTLplan is a forward-chaining planner, which basically adds one action at a time, introducing conditional branches whenever several observations are possible (like after the smell action in Figure 3), and keeping the search space down by the use of strategic knowledge formulated as logical formulae. It searches for a plan with a probability of success above a given threshold, and returns such a plan as soon as it is found. As mentioned, certain situations may arise where recovery planning may be necessary. This may require the robot to gather more information about the different perceived objects in order to resolve ambiguities or reacquire existing objects.

III. Human-Computer Interface

In addition to the computations mentioned in the previous section for control, perception and autonomy, the system also has a number of processes for displaying the internal state of the robot as well as its current model of the external world. This model of the external world includes locally perceived objects and a gridmap of the environment built from sensor data. This is shown in Figure 4. In the LPS window different sensing modules can be selected, as well as tracing the movements of the robot. Furthermore, a live feed of the images viewed from the CCD camera is available and the robot can be either controlled autonomously or through the cursor keys on the keyboard.

A novel feature of the interface is an olfactory component, which translates an electronic perception of odours into a human domain. The interface also elaborates on commands from a human into specific planning and sensing tasks on the robot. A snapshot of this interface is shown in Figure 4. The interface allows for different odour repositories to be loaded and stored, in the right most column. Each odour contains a signature, label L_o, and a hidden field which represents its stored anchor to an object in the real world, if available[2]. The classification process described in Section II-A is executed in the middle bar of the window. The process can be edited to adapt for cluster size, initialization algorithms, and logging outputs. As default the cluster size is set to the number of different names in the current repository. In the lower view a symbolic output of the clustering process is shown, where the clusters names are obtained from the contingency analysis. This gives explicit information about how the electronic nose categorizes the odours. Finally, a visualization of the clusters centers as well as the data points is shown in the upper left window (in 2-dimensions). When a particular point on the data space is selected by clicking the mouse, a plan to visit the anchor that refers to the selected point is generated and the plan is executed. For example, objects which correspond to points located in clusters with high uncertainty (i.e. low density) can be re-sampled by the robot when selected by the user.

IV. Experimental Validation

Keeping with our long term ambition to have a robot that works online in an environment for an extended period of time, our experiments statistically validate the performance of such a system. Two experiments examine the ability to acquires new knowledge of odours online and its affect on classification performance. Two additional experiments show the execution of tasks that a human can request from the robot.

[2]Some odours signatures may be unanchored as shown in Section IV-B

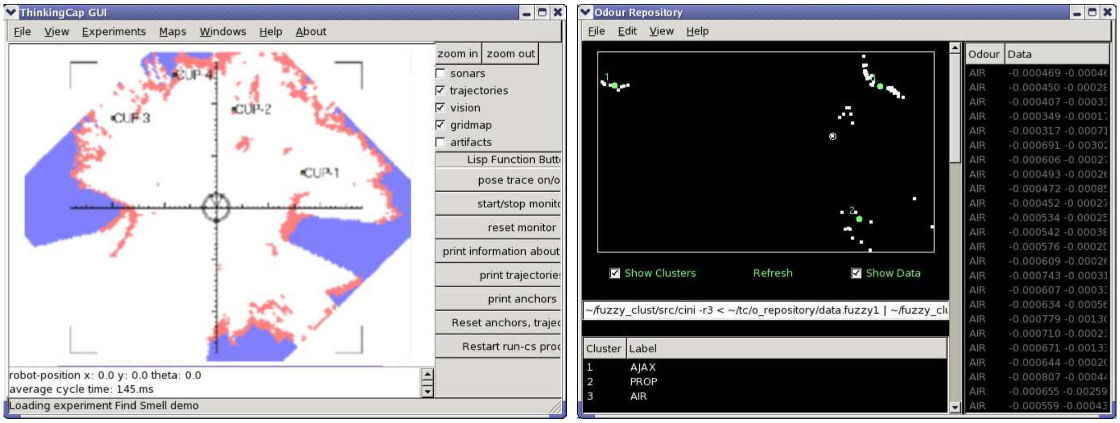

Fig. 4. *(Left)* The LPS window shows the detected objects, the robot is located in the center of the space. Grey areas in the space are regions that are unexplored. Objects are represented by their visual trajectories are placed within the LPS. In this screen, the robot identifies 4 cups. *(Right)* The olfactory interface shows clusters and data points for a loaded repository. Certain points can be interactively selected to generate new plans for perceptual actions. It also shows the symbolic representation preserving the electronic perception of odours use to classify new odours.

A. Acquisition of Training Data

Before any odours can be identified, training samples need to be collected. Bluntly said, this can be a tedious and monotonous process, especially in olfaction when the sampling of an odour including baseline, and purging cycles can take up to 2 minutes. An autonomous robot that could collect samples would be an attractive feature. However, often expert knowledge is needed for to evaluate the results, identify anomalous readings, and in some cases actively select the next training sample. In the interface presented in the previous section, functionalities are available so that a human and the robot can cooperate to collect training samples. The question is whether classification improves using this process and can a robot conduct this process without the supervision of the human. To address these issues, an odour repository was built using odour samples collected over a period of 2 days. The experiment was repeated twice, in the first case a human user interacting with the system determined which samples to take and when. In the second case, the robot autonomously determines which samples to acquire based on a measure of the uncertainty in the existing training set. This measure, is based on the results obtained from the classification algorithm, and is calculated as follows:

$$\mu_k = max(u_{ik}) \tag{3}$$

for $i = 0$ to n clusters. The next candidate sample is the k^{th} sample with the lowest μ_k (i.e. least membership to all clusters). The signatures of at least 4 different odours were contained in the repository, with $n = 4$ as a given parameter. Through each sampling cycle, the validity measures were evaluated where a total of 60 sampling cycles were performed. Each sampling cycle contains 5 samples of each odour. During the experiment, the coefficient of uncertainty averaged 0.97 ± 0.03, meaning that the categories from the human domain maintained correlation with the clusters in the electronic domain.

Fig. 5. Clustering performance shown in the compactness and separation validity measure during the training of 4 different odour. Clusters improve quality both when the robot is guided by the human, and when it determines the next best sample autonomously.

Figure 5 shows the graphs of the Compactness and Separation measure mentioned in Section II-A. Smaller values of Compactness and Separation measure indicates good clusters (dense, compact and well-separated). In both cases, as the number of samples increase with selective training, the global properties of the clusters improve. This result indicates that a robot can be used to collect the training samples for odours, and that this process can be adaptive based on the properties of the sensor data.

B. Online Learning of Odours

An important task is to be able to acquire new knowledge about objects as they become available to the system. This is important if new odours are encountered during the course of an experiment. To illustrate the robot's on-line ability to acquire new odours and then use that odour property to recognize new objects, we perform a search scenario where a human and the robot cooperate to find an object of a particular odour character. Here, we consider that the human is at a disadvantage and cannot describe the odour character, nor has the robot encounter the odour previously. So to learn the new odour, the human presents to the robot an artifact which contains the odour property, in this case a piece of cloth sprayed in 3-hexanol, Figure 6 (*Left*). The first task of the robot is to sample the artifact and store the odour property. The next task for the robot is to enter a room with many potential objects, determine the possible candidates with the same odour property, and then generate a plan to visit each candidate until a match is found. An example of the trace of the robot's movements is shown in Figure 6(*Right*).

Of the total 27 runs where a matching object was present, the robot succeeded to acquire it in 74% of the trials. Through each run, the number and position of candidate objects varied, as well as the starting position. Failures were caused by the detection of mirage objects (18.5%), misclassification of of the target object's odour (11.1%) and misdetection of present objects (11.1%). Note that in certain failed trials, candidate objects were missed *and* phantom objects were detected resulting in the overlap in percentage errors.

C. Object Recognition using Olfaction

In this experiment, the human requests the robot to find an object based on the olfactory property of that object. The objects are first seen using vision and then distinguished using olfaction. The planner, described in Section II-E, creates a plan that determines how each of the objects are visited and when the smell action is executed.

The following experiment shows the results when several visually similar objects are placed in a room in a variety of configurations. Up to 5 different objects are tested, where the only discriminatory feature between them is their odour character. The results from the runs are shown in Table I.

TABLE I
EXPERIMENTAL RESULTS FROM AMBIGUOUS CASES

Number of Odours	Trials	Olfactory Failures (Vision, Odometry)	Olfactory Failures (Classification)
2	11	18%	0 %
3	15	20%	0 %
4	21	19%	4.7%
5	25	16%	8%

These results show that as the number of odours increase the greater chance of misclassification, this result is expected and unfortunately is due to the sensing limitations of the e-nose. However, the majority of olfactory failures are a direct result of visual and odometry errors. This demonstrates that in cases of odour classification it is important that the robot be correctly placed before taking smelling actions. The next experiment, explicitly addresses this problem.

D. Thorough Object Inspection

In this experiment, the human requests the robot to thoroughly inspect an object. Thorough inspection may be required to reduce the uncertainty in an odour classification. Sources of uncertainty may be due to incorrect positioning to take an odour sample, localization error which places the robot too far from the object, or an uneven distribution of the odour around objects due to awkward shape, or air currents. Consider a situation where a large black suitcase may contain traces of a "dangerous" substance. Since it is important to reduce the probability of a false negatives i.e. claiming the suitcase is benign when it is not, the planner considers that an object may need to be sampled repeatedly and from several sides, thereby increasing the probability that at least one attempt is successful. On the other hand, the electronic nose rarely detects phantom odours. In other words, there are few occurrences in which the sensors will react when no odour is present. Consequently, the planner's model currently does not take the possibility of false positives into account.[3]

The robot's planner was equipped with a domain model including an action for approaching and docking to a suitcase or other objects from a given angle (0, 90, 180 or 270 degrees), with a certain probability that the robot is inadequately positioned after that action. There was also a smell action, which, if the robot was adequately positioned, would detect any "dangerous" substance. In reality a cleaning agent that was lightly sprayed on one side of the suitcase was used as a stand-in for the explosive TNT. The side of the suitcase that contains the odour is called the emission angle. From the planner's perspective, each sampling attempt has a certain probability p to fail ($\approx$ 20%), and therefore the accumulated probability of n consequitive failures ($1 \leq n \leq 4$) is p^n.

The results of 40 trials showed that the robot was able to correctly assess the suitcase 81.16% of the trials when an odour source was present, and 100% when no odour was present. Most often the optimum detection of an odour would occur when the robot approached the object from the emission side. However, of the 81.16% successes, about 37% found an odour at when the robot was placed 90° from the emission point and 18% when the robot was placed 180° (on the opposite side of the suitcase).

The benefit of calling additional smelling actions can be seen in the order in which movements were made. Suppose detection can be described in three terms: early detection, punctual detection and late detection. When the emission

[3]It is however well within the capacity of the planner to do so, and it might be desirable in situations where e.g. there are other suspicious suitcases very close the the one being investigated, and the odour might come from one of those.

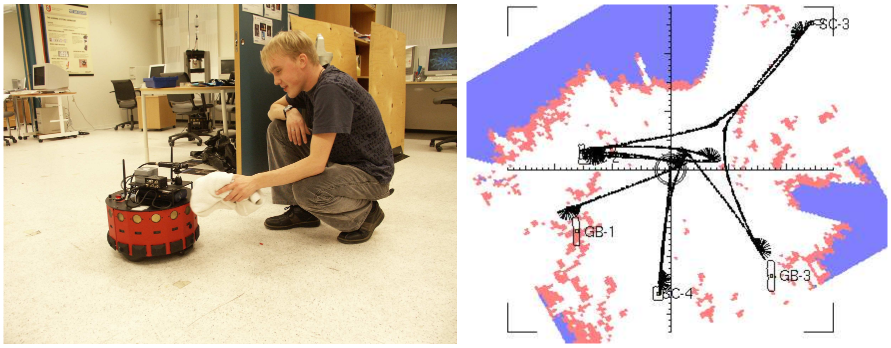

Fig. 6. *(Left)* Human presenting an odourous object *(Right)* The trace of the robot's movement during the search for an object of a particular odour in a room containing gasbottles, suitcases, boxes, and other objects.

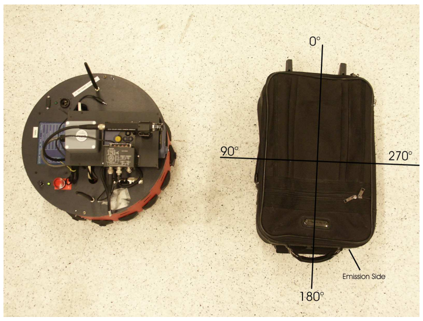

Fig. 7. The robot approaching a suspicious suitcase for careful inspection.

side is located at 180°, and the robot visits and "sniffs" the suitcase in the following order: (90°, 180°, 270°, 0°), as shown in Figure 7. Early detection would occur if the object was correctly assessed at 90°, late detection at (270° or 0°) and punctual detection at 180°. In the experimental runs, early detection occurs in about 29% of the successful trials and late detection occurs in approximately 23%. Therefore, planning for additional smelling actions in this case was able to increase the total success rate by preventing the robot from drawing premature conclusions about the nature of the package before careful inspection.

V. Conclusion and Future Work

The development that we expect, is that electronic olfaction will grow to be a standard equipment on many mobile robotic systems, from field robots to humanoids. Much in the same way as the computer vision community has developed different techniques and algorithms suited for mobile robots, an equal development is needed in electronic olfaction for smelling robots. This development should focus on all the aspects of artificial olfaction. In this work we considered a few of these aspects, mainly, the ability to discriminate objects based on odour property, the exploitation of the mobility of the robot to improve classification performance, the need for a human-computer interface in order to interact with a new perceptual system. Future work will be dedicated the to the refinement of the methods presented here, namely an improvement of the role of the planner to provide optimal navigation when visiting several objects, better reasoning about perceptions in order to refine the training process, and working in an online environment for extensive periods of time in a realistic monitoring context.

References

[1] S. Coradeschi and A. Saffiotti. An introduction to the anchoring problem. *Robotics and Autonomous Systems*, 43(2-3):85–96, 2003.

[2] D. Dubois. Categories as acts of meaning: The case of categories in olfaction and audition. *Cognitive Science Quaterly*, 1:35–68, 2000.

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

[4] I. Gath and A. Geva. Unsupervised optimal fuzzy clustering. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 11(7):773–781, 1989.

[5] Cyranose Sciences Inc. The cyranose 320 electronic nose, 2000. User's Manual Revision D.

[6] H. Ishida, T. Nakamoto, and T. Moriizumi. Remote sensing of gas/odor source location and concentration distribution using mobile system. *Sensors and Actuators B*, 49:52–57, 1998.

[7] A. Lilienthal and T. Duckett. Creating gas concentration gridmaps with a mobile robot. In *Proceedings of the 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2003)*. IEEE Press, 2003.

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

[9] A. Loutfi and S. Coradeschi. Forming odour categories using an electronic nose. In *Proc. of European Conference in Artificial Intelligence (ECAI 2004)*, pages 119–124, 2004.

[10] A. Loutfi, S.Coradeschi, L. Karlsson, and M. Broxvall. Putting olfaction into action: Using an electronic nose on a multi-sensing mobile robot, 2004. To appear in IEEE International Conference on Intelligent Robots and Systems (IROS 2004).

[11] R. Russell, D. Thiel, and A. Mackay-Sim. Sensing odour trails for mobile robot navigation. In *In IEEE Int. Conf. Robotics and Automation (ICRA 1994)*, pages 2672–2677, 1994.

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

[13] M. Wandel, A. Lilienthal, T. Duckett, U. Weimar, and A. Zell. Gas distribution in unventilated indoor environments inspected by a mobile robot. In *Proceedings of the IEEE International Conference on Advanced Robotics (ICAR 2003)*, Coimbra, Portugal, 2003.

[14] X. Xie and G. Beni. A validity measure for fuzzy clustering. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 13(8):841–847, 1991.