An Ecological Approach to Odour Recognition in Intelligent Environments

Mathias Broxvall, Silvia Coradeschi, Amy Loutfi and Alessandro Saffiotti
Center for Applied Autonomous Sensor Systems
Örebro University, Örebro, Sweden 70182
www.aass.oru.se

***Abstract*— We present a new approach for odour detection and recognition based on a so-called PEIS-Ecology: a network of gas sensors and a mobile robot are integrated in an intelligent environment. The environment can provide information regarding the location of potential odour sources, which is then relayed to a mobile robot equipped with an electronic nose. The robot can then perform a more thorough analysis of the odour character. This is a novel approach which alleviates some the challenges in mobile olfaction techniques by single and embedded mobile robots. The environment also provides contextual information which can be used to constrain the learning of odours, which is shown to improve classification performance.**

I. INTRODUCTION

Mobile robots equipped with artificial olfaction capabilities have a great application potential in a wide spectrum of domains. An olfaction-capable robot could detect hazardous or anomalous situations in private, public, or industrial environments; it could help in risk assessment and victim search during rescue operations; it could be used for medical monitoring in a home or retirement house; and so on.

In spite of this great potential, the integration of olfaction in a mobile robot is hindered by a number of technical difficulties. First, current devices for artificial olfaction ("electronic noses", or e-noses) are only able to perform reliable odour detection and classification when the concentration of odour molecules in the air is high. Thus, an interesting odour event would be missed unless the e-nose is very close to the odour source. Second, today's e-noses typically have a very slow reaction time. This, together with the complex nature of odour propagation, makes it very difficult for a mobile robot to correctly localize and approach an odour source. Finally, odour classification can only be done reliably if the set of classes is very small. This limits the applicability of current e-noses to cases in which we have some *a-priori* information of the type of substance which is being smelled.

Because of these limitations, the possibility of having a general olfaction-enabled robot in our homes, which can detect and classify abnormal events everywhere in the home, is not realistic today. In this paper, we explore a different road to approach this problem. Instead of assuming an all-mighty olfactory robot embedded in a passive environment, we assume a smart environment which contains a number of sensors and tagged objects, and which can provide the robot with the needed information. This approach is an instance of the general concept of PEIS-Ecology, which integrates autonomous robotics and intelligent environments. This general concept and its implementation are detailed in a companion paper [1] and in [10].

In the specific instance explored in this paper, we assume that the environment contains a number of very simple e-noses placed at critical locations (e.g., inside the refridgerator, or near the cooker). These simple devices can detect an abnormal gas concentration, but they are unable to classify the type of odour. We also assume that objects in the environment (e.g., goods in the refridgerator) can have tags attached, which contain information about the object itself.[1] When a simple e-nose detects an alarm, its location is sent to a mobile robot equipped with a sophisticated e-nose. The robot navigates to that place, and smells the different objects there. The information stored in the object's tags provides a context to restrict the classification problem.

In the next two sections, we discuss in more detail the current situation in electronic olfaction, and the PEIS-Ecology approach. We then describe an experimental system for odour recognition in a PEIS-Ecology framework, and show a few experiments that demonstrate the viability of our approach.

II. CURRENT SITUATION IN ELECTRONIC OLFACTION

The goal of electronic olfaction is to mimic the biological sense of smell using artificial sensors. The majority of the applications of electronic olfaction use an "electronic nose" which consists of two components: an array of gas sensors with partial or overlapping sensitivities and a pattern recognition component capable of discriminating between simple and complex odours. Other applications, mainly those in mobile robotics, may use simplified version of an electronic nose consisting of only a few gas sensors that do not necessarily recognize or classify an odour but rather react to the presence of a specific and a priori known component in the environment [9]. With that said, the current research interest in the area of electronic olfaction can be divided into two domains. The first domain deals with static applications whose focus is on the discrimination and identification of odour components. The second domain focuses on the mobile applications whose focus is on exploring an environment and finding an odour source. Both these domains have made significant progress, however, a number of challenges still remain.

[1] This assumption is not unrealistic: for example, many consumer goods already have an RFID tag attached to them.

In the field of odour identification and recognition, electronic noses are best used in the context of a specified domain. For instance, successful experiments have been conducted in the assessment of the quality of meat, wines, dairy products, eye infections and air quality [7]. In all these experiments the amount of different odour compounds in the training set typically range from 3 to 10. This is due to the current limitations in the gas sensing selectivities. As the number of different odours in the training set increases the probability to distinguish between these odours decreases.

Meanwhile the field of mobile olfaction has mainly neglected the classification problem in olfaction and focused on finding the source of just one and a priori known smell (e.g. cup of ethanol) [5]. Often biologically inspired search techniques or information about air distribution are used in combination with measure of the concentration gradient from the gas sensors [3]. The main challenges facing navigation by smell are due in part to the current limitation in the gas sensing technology and also to the stochastic nature of air turbulence in real environments. In order for gas sensors to react to an odour, the molecules of a particular compound need to come into contact with the surface of the sensor, which requires a large amount of the odourous molecule to have distributed in the environment. For the great majority of odours the necessary concentration levels for detection are often present a few decimeters from the odour source. Ethanol or alcohols which are often used for experiments exhibits exceptional distribution properties and cannot therefore be considered a typical example of odour distribution. Although odour molecules can be carried by air currents, the turbulent nature of air currents in real environments make it difficult to determine the source since often pockets of high concentration can be detected far from the source. Finally, the slow recovery time of the sensors require slow movement of a mobile robot which makes searching any large area time consuming and impractical.

III. The PEIS-Ecology Approach

The concept of PEIS-Ecology [10] puts together insights from the fields of autonomous robotics and ambient intelligence [4] to generate a radically new approach to building assistive, personal, and service robots. The main constituent of a PEIS-Ecology is a *physically embedded intelligent system*, or PEIS. This is any computerized system interacting with the environment through sensors and/or actuators and including some degree of "intelligence". A PEIS can be as simple as a smart toaster and as complex as a humanoid robot.

Individual PEIS in a PEIS-Ecology can co-operate based on the notion of linking functional components: each PEIS can use functionalities from other PEIS in the ecology in order to compensate or to complement its own. The power of the PEIS-Ecology does not come from the individual power of its constituent PEIS, but it emerges from their ability to interact and cooperate. As an example, a robot which needs to grasp a bottle would not use its sensors to detect its position, shape, and weight — a task which proved to be surprisingly difficult in years of robotic research. Instead, the bottle itself, enriched with a micro-PEIS, would hold this information and communicate it to the robot. More details of the concepts of PEIS-Ecology can be found in the companion paper [1].

Fig. 1. An example of olfaction in a PEIS-Ecology. The gas sensors provide detection and location of candidate odour alarms, the objects provide classification context.

The PEIS-Ecology approach can help us to address the challenges of robot olfaction discussed above. To detect distant odour, the olfactory robot could rely on the pervasive sensing functionality provided by a network of simple PEIS, each including a simple gas sensor, which are distributed at tactical locations in the environment. The same network of PEIS would provide the mobile robot (itself a PEIS) with location information about the potential odour source. When an odour alarm is generated, then the mobile robot, equipped with a much more sophisticated e-nose component, can move to that location and investigate the nature of the odour. The PEIS-Ecology approach can also address the problem of providing domain information for classification. Each tagged object that needs to be investigated can be seen as a PEIS in the ecology, and provide its identity to the mobile robot.[2] This identity can be used as a context to load the correct training set from a large odour repository, thereby increasing the probability of correct classification. This PEIS-Ecology is schematically illustrated in Fig. 1.

Compared to the tradition in electronic olfaction, the PEIS-Ecology approach includes two interesting innovative aspects. First, this approach reverses the tradition in electronic olfaction by placing the recognition and identification on a mobile platform and the simple gas sensor in static locations in the environment. This is a cost effective strategy considering that the simple sensors are cheap (5-20 USD per sensor) and small and can be easily distributed, whereas electronic nose devices are often quite costly. Second, this approach allows the introduction of context in olfaction. To the authors' knowledge, this is an aspect of odour discrimination that has not been previously explored.

[2] In the actual experiment described below, the refridgerator contains a RFID tag reader that can determine the identity of the (tagged) objects in it. Hence, the refridgerator, not the objects, is a PEIS.

IV. A PEIS-ECOLOGY FOR OLFACTION

We now proceed to describe a concrete system that applies the PEIS-Ecology framework to the problem of mobile olfaction. The resulting system is able to perform detection and identification of a variety of different odours in a large environment. In this section, we describe: the experimental platform used to implement and test PEIS-Ecologies; the specific PEIS-components used in our experiment; and the way in which they cooperate to solve our specific problem.

A. The PEIS*-platform*

We have constructed a reference implementation of a PEIS-Ecology, complete with olfactory capabilities, as well as built a demonstrator environment in which these tests can be performed.

The basic building blocks of our PEIS-Ecology is a "PEIS-component", that is, a logical process that implements an individual functionality inside a PEIS. In order for these PEIS's to cooperate, the PEIS-platform implements a distributed communication mechanism which allows the addition and removal of components from the ecology as needed. The main communication and synchronization mechanism is a distributed tuple-space enriched with an event mechanism. This is a generative communication method, which simplifies the dynamic reconfiguration of the PEIS-Ecology.

For the actual implementation of the software components used we developed a middle-ware called the PEIS-kernel, which provides these communication methods as well as perform other important services like network discovery and routing of messages between PEIS lacking direct means of communications [1].

In addition, we have built a physical test-bed facility, called the PEIS-home, which looks like a typical bachelor apartment of about $25m^2$. It consists of a living room, a bedroom and a small kitchen. The PEIS-Home is equipped with communication and computation infrastructure, and with a number of sensors. Fig.**??** shows a few snapshots of the home.

Fig. 2. *(Upper Left)* A rough sketch of the PEIS-Home. *(Upper Right)* A picture of the control deck. *(Lower Left)* A picture of the kitchen with a fridge located under the workbench. *(Lower Right)* A picture of the common room.

Fig. 3. *(Left)* The PEIS fridge showing the placement of the two gas sensors and the products to be sampled *(Right)* The door mechanism which opens and closes the fridge door.

B. The used PEIS

The refrigerator PEIS*:* The refrigerator PEIS consists of a apartment sized refrigerator ($51 \times 53 \times 81$ cm), two simple TGS 2600 Figaro gas sensors (sensitive to air contaminants) and a laptop computer with data acquisition technology [2]. Under current experimental conditions the refrigerator is turned off and the inside temperature is approximately 20.5°C. Milk products used during the experiments are placed in the lower shelf of the refrigerator as shown in Fig. 2. One sensor is placed in the shelf door (marked Sensor A in the figure) and the other sensor is placed on the roof of the refrigerator (Sensor B).

Sampling both sensors at a rate of 4 Hz an alarm is triggered when the concentration gradient exceeds a predefined threshold ε for both sensors.

The PEIS-fridge also incorporates an RFID tag reader (Texas Instruments IDISC MR 100). Each object stored in the refrigerator has an RFID tag attached, which stores the object's unique ID together with other optional properties (not used in this experiment). The PEIS-fridge inspects its contents and publishes it on the tuple-space at a fixed rate.

All the communications between the refrigerator PEIS and the other PEIS in the ecology happens by exchanging tuples through the distributed tuple-space. In particular, the refrigerator produces tuples of type `SmellAlarm`, containing the gas sensor data, and of type `Contents`, containing the ID's of all the objects inside the refrigerator. It responds to tuples of type `DoorCommand`, carrying requests to open or close the door.

Pippi: the mobile robot PEIS*::* The second PEIS in our ecology is an iRobot's Magellan Pro indoor research robot called Pippi. In addition to the usual sensors, Pippi is equipped with a CCD color camera and with a Cyranose 320$^{\text{TM}}$ elec-

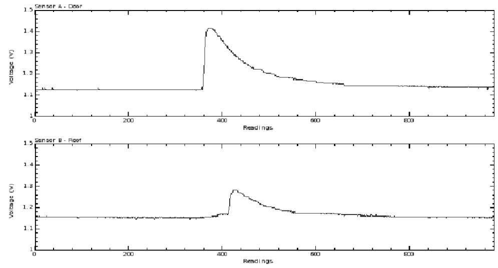

Fig. 4. Response from the two TGS 2600 sensors when an alarm is triggered.

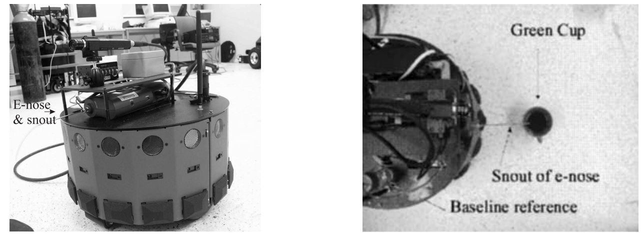

Fig. 5. *(Left)* A view of Pippi with the Cyranose 320 mounted underneath the camera. *(Right)* Top view of Pippi when sampling from a cup.

tronic nose used to identify and discriminate between odours. The Cyranose 320^{TM} electronic nose is a contained unit which relies on a 32-channel carbon-black polymer composite chemiresistor array for odour sampling. The sensors on Pippi are mounted in such a way that line of sight and line of "smell" are coordinated. This is done by placing the e-nose below the camera and odour samples are drawn from a uni-directional air-flow at the front of the robot, as shown in Fig. 4. The detection of an odour is therefore maximized when the robot is directly facing the source. and with a wireless network.

When sampling an odour using the Cyranose 320^{TM} a three phase sampling technique is used. The first phase, is a baseline collection where the sensors are exposed to a reference gas. The reference gas is contained in a 50 ml vile and consists of clean air that was prepared and collected in a ventilation cupboard. The vile is placed and fixed onto the mobile robot as shown in Fig. 4 (right). The baseline is collected for 30 seconds. The second phase is a sampling phase where odour samples are drawn into the e-nose unit and come in contact with the sensing array. The odour is sampled for 60 seconds. Finally, the third phase is a recovery where the sensing unit is purged of any reminiscent gases, this phase typically lasts for 60 seconds however, can be extended if the sensors do not exhibit a full recovery (i.e., return to approximate baseline values).

Sensor data pre-processing involves: (1) a fractional baseline manipulation, where the baseline is subtracted and then divided from the sensor response; (2) Local vector normalization, where the feature vector of each individual sniff is divided by its norm and lies on a hypersphere of unit radius; (3) Global auto-scaling in which the distribution of values for each sensor across the entire database is set to have zero mean and unit standard deviation. Identification of new odours is performed using a classification component, which uses the raw sensor data provided by the nose together with domain information and provides classified data. The feasibility of using an electronic nose on a mobile robot with the setup described above has been previously validated in [6].

On-board Pippi also runs an instance of the Thinking Cap, an architecture for autonomous robot control based on fuzzy logic [11], and an instance of the player program [8], which provides a low-level interface between the robot's sensors and actuators and the PEIS-Ecology's tuple-space.

Pippi responds to tuples providing commands and requests. These include tuples of type `Goal`, providing navigation goals, and of type `Smell`, providing smell commands. Pippi produces, among others, tuples that indicate the state of the navigation and the olfactory classification results.

The Home Security Monitor PEIS*:* This is a PEIS which consists of a static computer which is connected to a set of web-cameras mounted on the ceiling. In addition to other monitoring tasks, not relevant here, this PEIS contains a high-level component responsible for detecting possible problems using the available olfactory resources: the simple gas sensors which are present in the refrigerator-PEIS and possibly in other devices, and the electronic nose on Pippi.

This component reacts to alarms published into the tuple-space by the simple e-nose components. When an alarm is received, it uses a simple planner to send Pippi to the location of the alarm, and to perform an olfactory analysis.

Other PEIS*:* In addition to the PEIS described above, the PEIS-Home contains several other PEIS that can participate in the ecology such as other mobile robots. These components are not relevant to our olfaction ecology, and they will not be described further in this paper.

C. PEIS-*Ecology configuration*

In our olfactory PEIS-Ecology, we have connected the above PEIS and their components as shown in Fig. 5. The PEIS-fridge provides olfactory alarms to the deliberator in the Home Security Monitor PEIS, and it provides information about the refrigerator's contents as context to the odour classifier inside Pippi. Pippi's navigation component (Thinking Cap) receives navigation goals from the Home Safety Monitor, and returns execution states and odour classifications. Additional gas sensors, not considered here, could of course be distributed in the environment at other locations in order to extend the monitored area.

In practice, the above configuration is generated automatically by the Home Security Monitor by informing each PEIS-component of which tuples it should use as input and output. In general, the configuration of a PEIS-Ecology can change dynamically depending on the context. For instance, when Pippi enters an area covered by the external localization system, it may switch to using the global localization functionality provided by that system instead of its own self-localization.

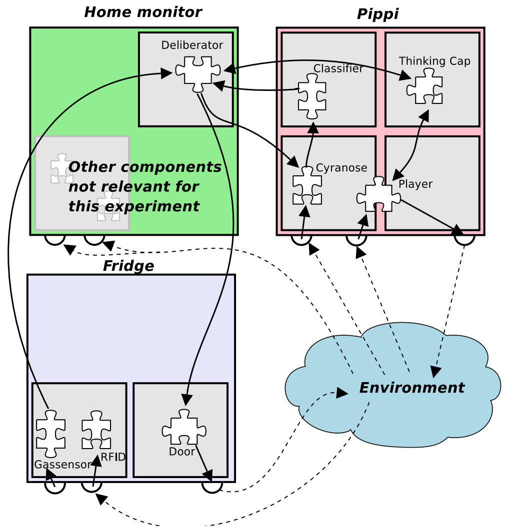

Fig. 6. The PEIS, PEIS-components, and connections used in our olfactory PEIS-Ecology. For simplicity the components not relevant to the olfaction experiment have been omitted.

V. Experimental Setup

In the first part of the paper, we have discussed two of the main obstacles to artificial mobile olfaction: (1) in order to obtain good classification of an odour it would be desirable to sample that odour close to the source [6], but general navigation by smell to approach the source is an unsolved problem; (2) good classification performance can usually be achieved if we strongly restrict the number of odour categories, but this requires to have some *a-priori* information about the odour source available. In our experiments, we undertake the following two appraoches to show that our PEIS-Ecology can help us to overcome these difficulties by: (1) using the location information provided by the distributed gas sensors to approach the odour source, and (2) using the identity information provided by the RFID-tagged objects to restrict the classification context.

All the experiments reported here have been performed using the PEIS-Ecology described in the previous section, and placing milk cartoons in the PEIS-fridge. Each experimental run consisted of the following:

- At start-up, Pippi started at a resting position in the living room.
- After some time, the gas sensors in the refrigerator triggered an alarm, and a tuple of type `fridgeAlarm` was placed in the tuple-space. This was notified to the deliberator component inside the Home Security Monitor PEIS, thus triggering the generation of an investigation plan.
- This plan was executed by first sending to the Thinking Cap component in Pippi the goal tuple `<B-Goal (at me refrigerator)>`.
- Pippi was dispatched to the refrigerator and when Pippi had moved in front of the refrigerator, thus achieving this goal, it announced success by publishing the corresponding tuple.
- The deliberator then activated the electronic nose by publishing another tuple, and waited for the classification results.
- If these results signified no problem, the deliberator asked Pippi to return to its starting position and resumed its monitoring activity.
- In all other cases, the deliberator would ask Pippi to move to the bedroom to alert the occupants of the PEIS-Home.

The focus of the experimental validations are directed towards the olfactory components, and the performance of the localization, deliberation and communication components are not given here. To evaluate how the contextual information affects classification, the training results of a classifier using two different repositories were compared. The first repository used is shown in the first column of Table I. This repository is not filtered according to the context and contains samples of a wide range of odours. The second repository which is a subset of the first contains samples that refer to the objects located in the refrigerator. The RFID reader in the refrigerator is able to detect which objects are present and publish their ID and information.

In each of the repositories, each record or "odour-print" represents an individual sample of an odour. Each sample consists of an odour signature or response vector $\vec{R_k}$ where each sensor's reaction is represented by:

$$R_{k,s} = \frac{R_{max,s} - R_{o,s}}{R_{o,s}} \qquad (1)$$

for s=1 to 32 sensors. $R_{max,s}$ is the maximum possible signal strength and $R_{o,s}$ is the calibrated baseline response for odour o.

VI. Experimental Results

A minimum of 10 samples were made for each odour used in each of the repositories (recall that each sample takes approx 2 minutes to collect). To validate context dependent training three different classification techniques were used in the classification module: K-means, KNN (with K=5 for both repositories), and canonical discriminant analysis. The results from the cross-validation are given in Table II. For KNN and K-means cross validation is performed by removing one exposure at a time. For canonical discriminant analysis, cross validation is performed by removing one exposure from each class. All the exposures are left out once during the cross validation procedure. The partial training set remaining and the removed exposures are scaled according to the pre-processing described in Sec. IV-B. For all three techniques, better classification performance was achieved when contextual information was used. Notice that classification performance is greater for the canonical discriminant analysis, this is most likely due to the fact that the canonical analysis is using Mahalonobis distance to determine inter-class distance relationships, whereas KNN and K-means use the Euclidean distance. The Mahalonobis

TABLE I
LIST OF SUBSTANCES IN EACH REPOSITORY

Repository A (No-Context)	Repository B (Context=Fridge)
Milk Quality-Good (14.5 °C)	Milk Quality-Good (14.5 °C)
Milk Quality-Good (19.0 °C)	Milk Quality-Good (19.0 °C)
Milk Quality-Poor (14.5 °C)	Milk Quality-Poor (14.5 °C)
Milk Quality-Poor (21 °C)	Milk Quality-Poor (21 °C)
Ethanol (C_2H_6O)	Plain Yogurt
Hexanal 98% ($C_6H_{12}O$)	Vanillin (C_8H_8O)
Linalool 98% ($C_{10}H_{18}O$)	
Ammonia (H_3N)	
Clean Air	
Plain Yogurt	
Vanillin (C_8H_8O)	
Octanol ($C_8H_{18}O$)	
Hexanoic Acid ($C_6H_{12}O_2$)	
3-Hexanol ($C_6H_{14}O$)	
Lavender	

TABLE II
CROSS-VALIDATION RESULTS FROM NON-CONTEXT AND CONTEXT DEPENDENT TRAINING

Classification Method	No Context (%)	Context (%)
K-means	79.3	92.5
KNN (5)	89.5	96.2
Canonical	92.5	94.8

distance is said to be superior measure since it takes the distribution of the points into account.

When using canonical discriminant analysis to classify new samples, the sample with the shortest Mahalonobis distance between its centroid and the sample in the canonical space is used to determine the class of the new sample. Table III gives the discrimination results for varying qualities of milk. The amount of unknown samples classified is given as well as the success of classification in percentage for the first repository (no-context) in Column 3 and for the contexted repository in Column 4. Since it was difficult to get fresh samples of milk to match the exact temperatures given in the training set, two temperature ranges were used. Milk at a cold temperature is considered to be less than 15°C and above 10°C and Room temperatures is above 15°C and below 20°C.

TABLE III
IDENTIFICATION OF UNKNOWN SAMPLES MILK QUALITY

Quality	Samples	Repository A	Repository B
Quality-Good (Cold)	20	60 %	90 %
Quality-Good (Room)	18	45 %	90 %
Quality-Bad (Cold)	20	70 %	85 %
Quality-Bad (Room)	20	85 %	100 %

Notice that the identification performance of Milk Quality-Good at Room temperature is low when using Repository A. The cause for this is a misclassification with other "mild" odours such as the clean air and the vanillin. However, by using context this classification result is greatly improved thus justifying approach (2) in Section V. Appraoch (1) has been implicitly justified since collection of odour samples would not have been possible without location information provided by the refrigerator.

VII. CONCLUSION

The main result of this paper is to show that we can use the concept of PEIS-Ecology to provide a new and viable solution to the problem of mobile olfaction. The PEIS-Ecology approach by-passes some of the difficult problems of mobile olfaction by exploiting the cooperation between the robot and an augmented environment. It should be noted that just using a network of artificial noses distributed everywhere in the environment would not be realistic: artificial noses able to correctly classify a reasonable range of odour are complex devices, which are still too massive and expensive to be distributed in the environment. By contrast, our approach relies on the combination of a network of simple, cheap and small sensors to detect a possible abnormal situation at a given location; and of a mobile robot equipped with a sophisticated artificial nose to move to that location and perform an in-depth analysis of the cause of the alarm.

ACKNOWLEDGMENT

This work has been supported by: ETRI (Electronics and Telecommunications Research Institute, Korea) through the project "Embedded Component Technology and Standardization for URC(2004-2008)" and Vetenskapsrådet.

REFERENCES

[1] M. Broxvall, M. Gritti, A. Saffiotti, B. Seo, and Y. Cho, "PEIS ecology: Integrating robots into smart environments," in *In Proc. IEEE Int. Conf. on Robotics and Automation (ICRA 06)*, 2006.
[2] "Figaro gas sensor: Technical reference," 1992, www.figarosensor.com.
[3] F. Grazzo, J. Basil, and J. Atenna, "Toward the convergence: robot and lobster perspectives of tracking odors to their source in the turbulent marine environment," in *In proceedings of the IEEE Intl. Symp. on Intelligent Control (ISIC)*, 1998, pp. 259–264.
[4] IST Advisory Group, "Ambient intelligence: from vision to reality," in *Ambient Intelligence*, G. Riva, F. Vatalaro, F. Davide, and M. A. niz, Eds. IOS Press, 2005.
[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)*, Seoul, South Korea, 2001, pp. 4005–4010.
[6] A. Loutfi, M. Broxvall, S. Coradeschi, and L. Karlsson, "Object recognition: A new application for smelling robots," *Robotics and Autonomous Systems*, vol. 52, pp. 272–289, 2005.
[7] T. Pearce, *Handbook of Machine Olfaction*. Wiley-VCH, 2003.
[8] Player/Stage Project, playerstage.sourceforge.net/.
[9] R. Russell, *Odour Detection by Mobile Robots*. London, UK: World Scientific Publishing Co, 1999.
[10] A. Saffiotti and M. Broxvall, "PEIS ecologies: Ambient intelligence meets autonomous robotics," in *Proc of the Int Conf on Smart Objects and Ambient Intelligence (sOc-EUSAI)*, Grenoble, France, 2005, pp. 275–280.
[11] A. Saffiotti, K. Konolige, and E. H. Ruspini, "A multivalued-logic approach to integrating planning and control," *Artificial Intelligence*, vol. 76, no. 1-2, pp. 481–526, 1995.