Putting Olfaction into Action: Anchoring Symbols to Sensor Data Using Olfaction and Planning

Amy Loutfi, Silvia Coradeschi, Lars Karlsson and Mathias Broxvall
Center for Applied Autonomous Sensor Systems
Örebro University
Örebro, Sweden 701-82
www.aass.oru.se

Abstract

Olfaction is a challenging new sensing modality for intelligent systems. With the emergence of electronic noses (e-noses) it is now possible to train a system to detect and recognise a range of different odours. In this work, we integrate the electronic nose on a multi-sensing mobile robotic platform. We plan for perceptual actions and examine when, how and where the e-nose should be activated. Finally, experiments are performed on a mobile robot equipped with an e-nose together with a variety of sensors and used for object detection.

1 Introduction

Mobile robots are becoming equipped with more diverse and numerous sensing modalities. Still, however, many of these modalities focus on obtaining the structural properties of an environment using for example, cameras, sonars and lasers. In object detection applications, objects may be discriminated more than just on their structural properties. One example is the property of smell. With the development of electronic olfaction and small electronic noses, gas sensors can be effectively integrated on mobile platforms. Furthermore such an integration could be of benefit to a number of robotics related applications for rescue robots, home-care robots and exploratory robots.

Electronic noses (e-noses) offer the potential to systematically quantify odours and this ability is attractive to many research and industrial applications. E-noses have also become commercially available, facilitating the integration of artificial olfaction in the AI community. Over the past two decades e-noses have been used in the detection of substances in the context of quality control in the food industry, detection of toxins, and discrimination of odours [Persaud and Dodd, 1982]. In this work we present a new kind of application where e-noses are added in the context of a larger sensing system, including vision, and planning. The e-nose is used for detecting objects where the odour characteristic is considered in a symbolic context as an additional property of the object. We focus in particular on the use of planning to determine the sensory actions needed to correctly collect the olfactory data.

Using an electronic nose in the context of object detection has several benefits. Firstly, descriptions of objects can include an odour and the respective recognition can be based on smell. For example, in the detection of “a cup of ethanol”, vision may serve to detect certain properties of the cup or even the colour of the substance but the actual substance of ethanol is best determined through chemical analysis. Another benefit emerges when two similar objects create an ambiguous case for the decision-making system. The property of smell can be used to disambiguate and distinguish between these objects. Even in non-ambiguous cases, acquiring the odour characteristic can be used to confirm the belief that the desired object has in fact been found. Finally, in the case of reacquiring an object, the odour property can be determined, stored, and reused in order to find the same object again.

A naïve approach to implementing odour recognition in the experiments mentioned above, 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 (2-5 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 is used instead as an active sensor that is explicitly called upon inside a complex decision making system.

The work in this paper presents an overview of a system that receives as input a symbolic description of an object and locates that object in a complex environment. The system has different sensing modalities available (in this case olfaction and vision) and should determine when to use these modalities in order to most effectively reach the goal. The system also reasons about its perceptions on a symbolic level. For this reason, a planner and an anchoring module are included within the system architecture. The purpose of the planner is to treat ambiguities, generate plans, call the execution of appropriate behaviours (that may include the command to smell an object) and determine when the goal has been reached. The anchoring component is important since it creates and maintains the connection between the symbols denoting objects at the symbolic level and the perceived physical objects [Coradeschi and Saffiotti, 2000].

The paper begins with a description of related works in Section 2. In Section 3, an overview of the system and its components is given. In Section 4, different scenarios such as disambiguating between objects and reacquiring objects are tested using vision together with olfaction on a robotic platform.

2 Related Work

To the knowledge of the authors, very little if not any attempts have been made to integrate an electronic nose on a multi-sensing mobile platform that specifically executes object recognition tasks. Although significant contributions have been made in the areas on odour source localization using mobile robots, it should be emphasized that this topic is outside the focus of this paper. Despite the absence of similar works however, there is well-developed research from both the electronic nose community and the AI community that facilitates the integration of odour recognition into intelligent systems.

In [Gardner and Bartlett, 1999] a general definition of an electronic nose includes both an array of gas sensors of partial selectivity and a respective pattern recognition process to detect simple or complex odours. Both a variety of sensing technologies (metal oxide semiconductor, conducting polymers, acoustic wave devices and fibre-optic sensors) as well as pattern recognition techniques have been applied in research, industrial and commercial domains. The most common of the pattern recognition used has been artificial neural networks, which are trained on odour categories. Although ANN's have provided good results in applications with limited amount of odours categories as shown in [Keller *et al.*, 1996], using black-box classification fails to address the problem of representing the knowledge of odour categories in order to classify a larger spectrum of odours. A study on the meaning of categorisation by Dubois [2000] has attributed the absence of fixed standards in odour classification to the lack of correlation between the chemical composition of an odour and the common name that refers to it. Consequently, chemical models such as Dravniek's [2000] character profiles or Amoore's [1965] primary odour tables, have not adequately represented in a generally accepted manner odour categorisations due to the large possibility of odour categories and descriptions. Thus, in recent years, there has been a movement to treat e-nose data in a more human-like manner by either relying on expert knowledge using fuzzy-based logic in tailored applications [Lazzerini *et al.*, 2001] or using explicit symbolic descriptions in order to relate odours categories to one another [**?**].

Aside from performing navigation by smell [Lilienthal *et al.*, 2001], little work has explored the possibility of using an electronic nose together with other sensors on robotic platforms. Some studies such as [Sundic *et al.*, 2000] have integrated e-noses on multi-sensing stationary platforms for the purpose of sensor fusion, but have not considered an application with intelligent or autonomous systems. Conversely, many intelligent or autonomous systems that consider planning for perceptual actions [Barrouil *et al.*, 1998; Kovacic *et al.*, 1998; Wasson *et al.*, 1998] and perceptual anchoring [Coradeschi and Saffiotti, 2003], have focussed primarily on vision-based sensors.

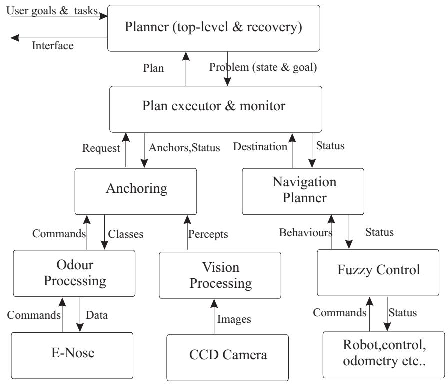

Figure 1: Overview of the robotic complete system which uses an anchoring module. On the left side of the arrows information is flowing downward and on the right side information is flowing upward.

3 An overview of the complete system

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. Further details regarding the other components of the system can be found in their respective references.

3.1 Planner

The planner, PTLplan, is a conditional progressive planner [Karlsson, 2001]. It has the capacity to:

- reason about incomplete and uncertain information: it might be unknown whether a cup contains ethanol or some other substance.
- reason about informative actions: by smelling the cup, one might determine whether it contains ethanol or something else.
- generate plans that contain conditional branches: if the cup contains ethanol, pick it up; otherwise, check the other cup.

PTLplan is used on two different levels: for constructing the plan for achieving the current main goal, and for constructing repair plans in case problems occur while the main plan is executed. The repair planning facility mainly deals with actions and plans for disambiguating an ambiguous situation, e.g. by finding out which of the two cups is the one that contains ethanol [**?**].

The planner functions by searching in a space of belief states. A belief state represents the agent's incomplete and uncertain knowledge about the world at some point in time. A belief 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 belief state, and in each belief state the robot will know more than before. To summarize, the planner searches for plans that maximize the probability of success, as well as mimimize the cost. Cost can be defined simply in terms of the number of steps, or by associating each action with a numerical cost.

3.2 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 repair planning facility is invoked.

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

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

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 cup-25" where "cup-25" is described as a "green cup".

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.

A more detailed description of the anchoring module and its functionalities can be found in [Coradeschi and Saffiotti, 2000].

3.4 Thinking Cap

In this system, execution monitoring on a mobile robot is controlled by a hybrid architecture evolved from [Saffiotti *et al.*, 1995] 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.

3.5 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. As long as our movements are limited this makes it easy for the system to reacquire objects based on their stored position, if however the objects move or if accumulation of odometry errors is large this might lead to reacquisition ambiguities which can only be resolved using other sensors.

3.6 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 an aluminium substrate. Thin films are deposited across the leads creating

Figure 2: A scenario where the robot discriminates four visually similar green cup objects based on their smell property

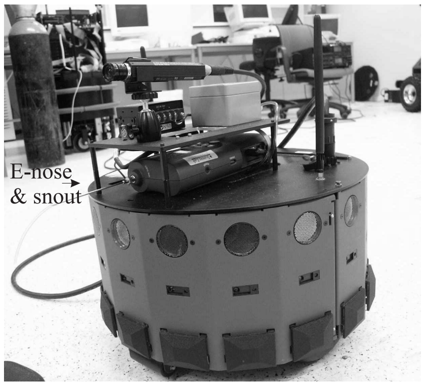

Figure 3: A closer view of the Pippi and the electronic nose.

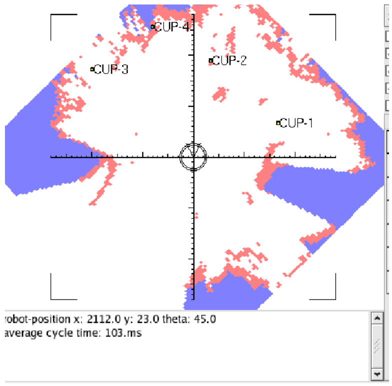

Figure 4: The local perceptual space of Pippi given 4 similar cups and obstacles. Pippi is located in the center of the space.

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 180 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 60 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. A total of 30 seconds are dedicated to purging the unit however the time for full recovery of the sensors may vary according to the odour being sampled.

The signals are gathered in a response vector where each sensor's reaction is represented by the fractional response of the reaction phase to the baseline. 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. With this type of data, the authors have used an unsupervised technique based on fuzzy clustering such as that presented in [**?**]. However, for the the experiments described below and in order to optimise computation time in a real-time environment, 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. Later work intends to explore the ability to output not only the class name but also a degree of membership, which could then be used in the planning process.

4 Experiments

The experimental setup consists of a Magellan Pro Research robot, called Pippi, equipped with a CCD camera, sonars, infrared sensors, 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 3. The nose has been previously trained on a number of different substances including those used in the first experiment.

4.1 Disambiguiting Objects

In this scenario, we consider a situation where Pippi is in a room where there are several different objects present as shown in Figure 4. On the floor of the room are visually identical green cups. Given the request to identify an object such as cup of ethanol, the planner in this case activates a Find functionality creating anchors for the matching objects. An example of the call to Find may look like:

(find c1 ((shape = cup) (smell = ethanol)))

The anchoring module examines the percepts sent from the vision module, and finds several percepts that match the indicated shape of c1. The percepts are considered as partially matching the description as the smell still remains to be perceived. The anchoring module classifies this as an ambiguous situation and creates several candidate anchors, one for each cup.

The repair planner is invoked, and considers the properties of the requested cup c1 and of the new percepts. It finds that c1 is expected to smell ethanol, but that there is no smell associated with either anchor (the case of two identical cups). It then automatically generates an initial state consisting of three hypotheses: that only anchor-1 smells ethanol and is the target for anchoring c1; the corresponding for the other anchors; and finally that no anchors smell ethanol. This initial state, and the goal to either have c1 anchored or to determine it cannot be anchored, is given as input to the planning algorithm, which then produces the following conditional plan:

```
((move-near anchor-1) (smell-obj anchor-1)
 (cond
    ((smell anchor-1 = ethanol) (anchor c1 anchor-1) :success)
    ((smell anchor-1 = not-ethanol) (move-near anchor-2)
     (smell-obj anchor-2)
     (cond
        ((smell anchor-2 = ethanol)
         (anchor c1 anchor-2) :success)
        ((smell anchor-2 = not-ethanol)
         (anchor c1 :fail) :success)))))
```

Note that each conditional branch ends with a decision to anchor c1 to one of the candidates, or to none (:fail).

The plan is executed and Pippi moves close enough to the object denoted by anchor-1 so that the snout of the nose is within reasonable sampling proximity (see Figure 4). The planner sends to the anchoring module a request to smell the object. The anchoring module receives the command and requests the sensing data from the e-nose. Now the e-nose invokes a 4-minute sampling procedure where the result is a classification based on the comparison with the trained values. The classification is sent and anchor-1 is updated. An indication that the smelling process has finished sampling is sent to the planner from the anchoring module. If the smell was found to be ethanol, the planner decides to associate the symbol c1 to anchor-1. If the object denoted by anchor-1 did not smell like ethanol then the planner would proceed to approach the second object. If neither of the objects denoted by the anchors return the desired classification, the planner registers that no anchors for c1 have been found.

A number of configurations of the above scenario were tested, where the number of cups were 2, 3, 4 or 5, each cup with a different content. The contents of the cups were one of the following: Ethanol, Hexanal, 3-Hexanol, Octanol, or Linalool. These substances are part of an ASTM atlas of odour descriptions [Dravnieks, 2000] whose characters are best described as alcoholic, sour, woody, oily and fragrant respectively.

Table 1 summarizes the results from different configurations involving different numbers of candidate objects. Note that in order 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 (accumulated odometry errors cause the anchoring module to lose track of an object). The table also provides information regarding the source of failures in the unsuccessful cases.

Analysis of the results shows that visual failures provoke olfactory failures. This is due to the fact that the 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 candidate objects increase, Pippi needs to move a longer distance and a larger odometry error is accumulated. There are cases in which the e-nose misclassifies 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.

Table 1: Experimental results from Disambiguating visually similar objects

No.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%

4.2 Reacquiring Objects

The purpose with this experiment is to show a different application of an electronic nose in which the e-nose is used to acquire the specific odour of an object. Assuming that the odour characteristic is a unique property of that object, this information can be used to reacquire the object again. In this scenario, the robot is in a room and a cup is located on the floor. Pippi is first given the task to find the cup and then to acquire its odour. Pippi first looks for the cup, finds it, and moves close to it. It then requests the e-nose to start sampling. In this case, however, the objective is not to recognise the smell but instead acquire it in order to use it for identifying the cup at future occasions. The e-nose samples the odour and stores the sensor signals as a new pattern within the training depository. A new name is generated for the odour. Finally, the anchoring module stores the information that this particular cup has the acquired smell.

The robot then wanders throughout the room. Meanwhile, an additional cup of similar shape and colour is added and the original cup is displaced so that it cannot be recognised by its position. The planner then requests the anchoring module to reacquire the original cup. Two possible candidates are found, the original cup and the new cup, and 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 second cup is checked. In our experiment, Pippi successfully reacquired the cup.

5 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 successful integration requires that the e-nose is explicitly called within the system at the appropriate occasion. Planning is essential for this task. The result is a system capable of using odour recognition to disambiguate between visually similar objects of different odour property and reacquire them at a later time.

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

References

[Amoore, 1965] J. Amoore. Psychophysics of odor. In *Cold Spring Harbor Symposia in Quantitative Biology*, volume 30, pages 623–637, 1965.

[Barrouil *et al.*, 1998] 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.

[Coradeschi and Saffiotti, 2000] S. Coradeschi and A. Saffiotti. Anchoring symbols to sensor data: preliminary report. In *Proc. of the 17th American Association for Artificial Intelligence Conf. (AAAI)*, pages 129–135, 2000.

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

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

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

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

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

[Keller *et al.*, 1996] P. Keller, L. Kangas, L. Liden, S. Hashem, and R. Kouzes. Electronic noses and their applications. In *World Congress on Neural Networks (WCNN)*, pages 928–931, San Diego, CA, USA, 1996.

[Kovacic *et al.*, 1998] 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.

[Lazzerini *et al.*, 2001] B. Lazzerini, A. Maggiore, and F. Marcelloni. Fros: a fuzzy logic-based recogniser of olfactory signals. *Pattern Recognition*, 34(11):2215–2226, 2001.

[Lilienthal *et al.*, 2001] 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.

[Persaud and Dodd, 1982] K. Persaud and G. Dodd. Analysis of discrimination mechanisms of the mammalian olfactory system using a model nose. *Nature*, 299:352–355, 1982.

[Saffiotti *et al.*, 1995] 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.

[Sundic *et al.*, 2000] T. Sundic, S. Marco, A. Perera, A. Pardo, J. Samitier, and P.Wide. Potato creams recognition from electronic nose and tongue signals: feature extraction/selection and r.b.f neural networks classifiers. In *Proc. of the IEEE 5th Seminar on Neural Network Applications in Electrical Engineering (NEUREL)*, pages 69–74, 2000.

[Wasson *et al.*, 1998] 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.