Grounded Symbolic Communication between Heterogeneous Cooperating Robots

DAVID JUNG†
Center for Engineering Science Advanced Research (CESAR), Computer Science and Mathematics Division (CSMD), Oak Ridge National Laboratory, Oak Ridge, TN 37831, USA
david.jung@pobox.com http://pobox.com/~david.jung

ALEXANDER ZELINSKY
Robotic Systems Laboratory (RSL), Department of Systems Engineering, Research School of Information Sciences and Engineering, The Australian National University, Canberra, ACT 0200, Australia
alex@syseng.anu.edu.au http://syseng.anu.edu.au/rsl/

Abstract. In this paper, we describe the implementation of a heterogeneous cooperative multi-robot system that was designed with a goal of engineering a grounded symbolic representation in a bottom-up fashion. The system comprises two autonomous mobile robots that perform cooperative cleaning. Experiments demonstrate successful purposive navigation, map building and the symbolic communication of locations in a behavior-based system. We also examine the perceived shortcomings of the system in detail and attempt to understand them in terms of contemporary knowledge of human representation and symbolic communication. From this understanding, we propose the *Adaptive Symbol Grounding Hypothesis* as a conception for how symbolic systems can be envisioned.

Keywords: cooperative robotics, heterogeneous systems, symbolic communication, symbol grounding, learning, representation, behavior-based, mobile robots

1. Introduction

The Behavior-based approach to robotics has proven that it is possible to build systems that can achieve tasks robustly, react in real-time and operate reliably. The sophistication of applications implemented ranges from simple reactivity to tasks involving topological map building and navigation. Conversely, the classical AI approach to robotics has attempted to construct symbolic representational systems based on token manipulation. There has been some success in this endeavor also. While more powerful, these systems are generally slow, brittle, unreliable and do not scale well – as their 'symbols' are ungrounded.

In this paper, we present an approach for engineering grounded symbolic communication between heterogeneous cooperating robots. It involves designing behavior that develops shared groundings between them. We demonstrate a situated, embodied, behavior-based multi-robot system that implements a cooperative cleaning task using two autonomous mobile robots. They develop shared groundings that allow them to ground a symbolic relationship between positions consistently. We show that this enables symbolic communication of locations between them.

The subsequent part of the paper critically examines the system and its limitations. The new understanding of the system we come to shows that our approach will not scale to complex symbolic systems. We argue that it is impossible for complex symbolic representational systems to be responsible for appropriate behavior in situated agents. We propose the *Adaptive Symbol Grounding Hypothesis* as a conception of how systems that communication symbolically can be envisioned.

Before presenting the system we have developed, the first section briefly discusses cooperation and

† This research was funded primarily by the The University of Wollongong and The Australian National University, and also in part by the Engineering Research Program, Office of Basic Enegry Sciences, of the U.S. Department of Energy, under contract DE-AC05-96OR22464 with Lockheed Martin Energy Research Corporation. The submitted manuscript has been authored by a contractor of the U.S. Government under the above contract. Accordingly, the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or allow others to do so, for U.S. Government purposes.

communication generally and looks at some instances of biological cooperation in particular. From this, we determine the necessary attributes of symbolic systems.

2. Cooperation and Communication

Cooperation and communication are closely tied. Communication is an inherent part of the agent interactions underlying cooperative behavior, whether implicit or explicit. If we are to implement a concrete cooperative task that requires symbolic level communication, we must first identify the relationship between communication and cooperative behavior. This section introduces a framework for classifying communication and uses it to examine some examples of cooperative behavior in biological systems. From this examination, we draw conclusions about the mechanisms necessary to support symbolic communication. These mechanisms are utilized in our implementation of the cooperative cleaning system, as described in the subsequent section.

It is important to realize that 'cooperation' is a word – a label for a human concept. In this case, the concept refers to a category of human and possibly animal behavior. It does *not* follow that this behavior is necessarily beneficial to the agents involved. Since evolution selects behavioral traits that promote the genes that encourage them, it will be beneficial to the genes but not necessarily the organism or species. Human cooperative behaviour, for example, is a conglomerate of various behavioural tendencies selected for different reasons (and influenced cultural knowledge). Because the design of cooperative robot systems is in a different context altogether, we need to understand which aspects are peculiar to biological systems.

2.1 Some Characteristics of Communication

Many authors have proposed classifications for the types of communication found in biological and artificial systems (e.g. see Arkin and Hobbs, 1992b; Balch and Arkin, 1994; Cao et al., 1995; Dudek et al., 1993; Kube and Zhang, 1994; Mataric, 1997a). Like any classification, these divide the continuous space of communication characteristics into discrete classes in a specific way – and hence are only useful within the context for which they are created. We find it necessary to introduce another classification here.

A communicative act is an interaction whereby a *signal* is generated by an *emitter* and 'interpreted' by a *receiver*. We view communication in terms of the following four characteristics.

- *Interaction distance* – This is the distance between the agents during the communicative interaction. It can range from direct physical contact, to visual range, hearing range, or long range.
- *Interaction simultaneity* – The period between the signal emission and reception. It can be immediate in the case of direct contact, or possibly a long time in the case of scent markers, for example.
- *Signaling explicitness* – This is an indication of the explicitness of the emitter's signaling behavior. The signaling may be a side effect of an existing behavior (implicit), or an existing behavior may have been modified slightly to enhance the signal through evolution or learning. The signaling may also be the result of sophisticated behavior that was specifically evolved, learnt, or in the case of a robot, designed, for it.
- *Sophistication of interpretation* – This can be applied to either the emitter or the receiver. It is an indication of the complexity of the interpretation process that gives meaning to the signal. For example, a chemical signal may invoke a relatively simple chain of chemical events in a receiving bacterium. It is possible that a signal has a very different meaning to the emitter and receiver – the signal may have no meaning at all to its emitter. Conversely, the process of interpretation of human language is the most complex example known.

2.2 Representation

It is not possible to measure the sophistication of the interpretive process by observing the signal alone. Access to the mechanics of the process within an agent would be necessary. Unfortunately, our current understanding of the brain mechanisms underlying communication in most animals is poor, at best. Therefore, our only approach is to examine the structure of the communicated signals. Luckily, there appears to be some correlation between the structural complexity of communicated signals and the sophistication of their interpretive processes. Insect mating calls are simple in structure and we posit a simple interpretive process. At the other end of the spectrum, human language has a complex structure and we consider its interpretation amongst the most sophisticated processes known. Bird song and human music are possible exceptions, as they are often complex in structure, yet have a relatively simple interpretation. This is due to other evolutionary selection pressures, since song also provides fitness information about its emitter to prospective mates and, in the case of birds, serves to distinguish between members of different species.

Science, through the discipline of linguistics, has learned much about the structure of the signals generated by humans that we call language (see Robins, 1997). We utilize a small part of that here by describing a conception of the observed structure of language. Using Deacon's terms we define three types of reference, or levels of representation: - *iconic*, *indexical* and *symbolic* (Deacon, 1997).

Iconic representation is by physical similarity to what it represents. The medium may be physically external to the agent – for example, as an orange disc painted on a cave wall may represent the sun. Alternatively, it may be part of the agent, such as some repeatable configuration of sensory neurons, or "*internal analog transforms of the projections of distal objects on our sensory surfaces*" (Shepard and Cooper, 1982).

Indexical reference represents a correlation or association between icons. All animals are capable of iconic and indexical representation to varying degrees. For example, an animal may learn to correlate the icon for smoke with that for fire. Hence, smoke will come to be an index for fire. Even insects probably have limited indexical capabilities. Empirical demonstrations are a mechanism for creating indexical references in others. Pointing, for example, creates an association between the icon for the physical item being indicated and the object of a sentence. The second part of this paper describes how we have used empirical demonstration to enable the communication of locations between robots.

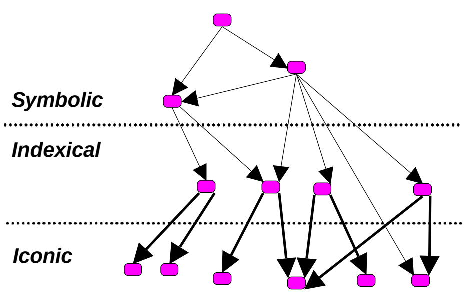

Figure 1 – Levels of representation

The third level of representation is symbolic. A symbol is a relationship between icons, indices and other symbols. It is the representation of a higher-level pattern underlying sets of relationships. It is hypothesized that language can be represented as a symbolic hierarchy (Newell and Simon, 1972). We will use the term *sub-symbolic* to refer to representations that need only iconic and indexical references for their interpretation. If the interpretation of a symbol requires following references that *all* eventually lead to icons, the symbol is said to be *grounded*. That is, the symbols at the top of Figure 1 ultimately refer to relationships between the icons at the bottom. A symbol's *grounding* is the set of icons, indices and other symbols necessary to interpret it.

The problems associated with trying to synthesize intelligence from ungrounded symbol systems – the classical AI approach – have been well documented in the literature. One such problem is termed the *frame problem* (see Ford and Hayes, 1991; Pylyshym, 1987). The importance of *situated* and *embodied* agents has been actively espoused by members of the behavior-based robotics community, in recognition of these problems, for many years (see Brooks, 1991; Mataric, 1992a; Pfeifer, 1995; Steels, 1996 for a selection). This hypothesis is called the *physical grounding hypothesis* (Brooks, 1990). Consequently, we adopted the behavior-based approach for our implementation.

2.3 Cooperative biological systems

In this subsection we describe five selected biological cooperative systems and classify the communication each employs using the scheme introduced above. From these, in the following subsection, we identify the necessary mechanisms for symbolic communication, which were transferred to the implementation of the cooperative multi-robot cleaning system.

2.3.1 Bacteria

Cooperation between simple organisms on earth is almost as old as life on earth itself. Over a billion years ago bacteria existed similar to contemporary bacteria recently observed to exhibit primitive cooperation. Biologists have long understood that bacteria live in colonies. Only recently has it become evident that most bacteria communicate using a number of sophisticated chemical signals and engage in altruistic behavior (Kaiser and Losick, 1993). For example, Mycobacteria assemble into multi-cellular structures known as fruiting bodies. These structures are assembled in a number of stages each mediated by different chemical signal systems. In these cases the bacteria emit and react to chemicals in a genetically determined way that evolved explicitly for cooperation. Hence, we classify the signaling as explicit. The interaction distance is moderate compared to the size of a bacterium, and the simultaneity is determined by the speed of chemical propagation. The mechanism for interpretation is necessarily simple in bacteria.

We can consider this *communication without meaning preservation* – the meaning of the signal is different for emitter and receiver. The emitter generates a signal without interpreting it at all (hence, it has no meaning to the emitter). The receiver interprets it iconically[1]. This type of communication has been implemented and studied in multi-robot systems. For example, Balch and Arkin have implemented a collective multi-robot system, both in simulation and with real robots, to investigate to what extent communication can

[1] The stereotypical way a particular chemical receptor on the bacteria's surface triggers a chain of chemical events within, is an icon for the presence of the external chemical signal.

increase their capabilities (Balch and Arkin, 1994). The tasks they implemented were based on eusocial insect tasks, such as forage, consume, and graze. One scheme employed was the explicit signaling of the emitter's state to the receiver. They showed that this improves performance, as we might expect. Specifically it provides the greatest benefit when the receiver cannot easily sense the emitter's state implicitly. This finding was also observed by Parker in the implementation of a puck moving task where each robot broadcast its state periodically (Parker, 1995); and by Kube and Zhang with their collective box pushing system (Kube and Zhang, 1994). The second part of this paper will demonstrate that the result also holds for our system.

2.3.2 Ants

Of the social insect societies, the most thoroughly studied are those of ants, termites, bees and wasps (Wilson, 1971; Wilson, 1975; Crespi and Choe, 1997). Ants display a large array of cooperative behaviors. For example, as described in detail by Pasteels *et al.* (Pasteels et al., 1987), upon discovering a new food source, a worker ant leaves a pheromone trail during its return to the nest. Recruited ants will follow this trail to the food source with some variation while laying their own pheromones down. Any chance variations that result in a shorter trail to the food will be reinforced at a slightly faster rate, as the traversal time back and forth is less. Hence, it has been shown that a near optimal shortest path is quickly established as an emergent consequence of simple trail following with random variation.

In this case, the interaction distance is local – the receiver senses the pheromone at the location it was emitted. As the signal persists in the environment for long periods, there may be significant delay between emission and reception. The signaling mechanism is likely to be explicit and the interpretation, while more complex than for bacteria, is still relatively simple. The ants also communicate by signaling directly from antennae to antennae.

Since both emitter and receiver can interpret the signal in the same way, we consider it *communication with meaning preservation*. The crucial element being that both agents *share the same grounding* for the signal. In this case, the grounding is probably genetically determined – through identical sensors and neural processes. This mechanism can also be applied to multi-robot systems. For example, if two robots shared identical sensors, they could simply signal their sensor values. This constitutes an iconic representation, and it is grounded directly in the environment for both robots identically. Nothing special needs to be done to ensure a shared grounding for the signal.

2.3.3 Wolves

A social mammal of the Canine family, wolves are carnivores that form packs with strict social hierarchies and mating systems (Stains, 1984). Wolves are territorial. Territory marking occurs through repeated urination on objects on the periphery of and within the territories. This is a communication scheme reminiscent of our ants and their chemical trails. Wolves also communicate with pheromones excreted via glands near the dorsal surface of the tail.

Wolves hunt in packs. During a pack hunt, individuals cooperate by closely observing the actions of each other and, in particular, the dominant male who directs the hunt to some extent. Each wolf knows all the pack members and can identify them individually, both visually and by smell. Communication can be directed at particular individuals and consists of a combination of specific postures and vocalizations. The interaction distance in this case is the visual or auditory range respectively, and the emission and reception is effectively simultaneous. The signals may be implicit, in the case of observing locomotory behavior, for example; or more explicit in the case of posturing, vocalizing and scent marking. It seems likely that the signals in each of these cases are interpreted similarly by the emitter and receiver.

Again, this is an instance of *communication with meaning preservation*. A significant difference is that the shared grounding enabling the uniform interpretation of some signals (e.g. vocalizations and postures) is not wholly genetically determined. Instead, a specific mechanism exists such that the grounding is partially learnt during development – in a social environment sufficiently similar to both that a shared meaning is ensured.

2.3.4 Non-human primates

Primates display sophisticated cooperative behavior. The majority of interactions involve passive observation of collaborators via visual and auditory cues, which are interpreted as actions and intentions. As Bond writes in reference to Vervet monkeys, "*They are acutely and sensitively aware of the status and identity of other monkeys, as well as their temperaments and current dispositional states*" (Bond, 1996). Higher primates are able to represent the internal goals, plans, dispositions and intentions of others and to construct collaborative plans jointly through acting socially (Cheney and Seyfarth, 1990). In this case, the interaction is simultaneous and occurs within visual or auditory range. The signaling is implicit but the sophistication of interpretation for the receiver is considerable. Some explicit posing and gesturing is also utilized, which is used to establish and control ongoing cooperative interactions.

As with the Wolves, we observe communication with meaning preservation through a shared grounding that is developed through a developmental process. In this case, the groundings are more sophisticated, as is the developmental process required to attain them.

2.3.5 Humans

In addition to the heritage of our primate ancestors, humans make extensive use of communication, both written and spoken, that is explicitly evolved or learnt. There is almost certainly some a priori physiological support for language learning in the developing human brain (Bruner, 1982). Humans cooperate in many and varied ways. We display a basic level of altruism toward all humans and sometimes animals. We enter into cooperative relationships – symbolic contracts – with mates, kin, friends, organizations, and societies whereby we exchange resources for mutual benefit. In many cases, we provide resources with no reward except the promise that the other party, by honoring the contract, will provide resources when we need them, if possible. We are able to keep track of all the transactions and the reliability with which others honor contracts (see Deacon, 1997 for a discussion).

Humans also use many types of signaling for communication. Like our primate cousins, we make extensive use of implicit communication, such as posturing (body language). We also use explicit gesturing – pointing, for example. Facial expressions are a form of explicit signaling that has evolved from existing expressions to enhance the signaling reliability and repertoire. Posturing, gesturing and speaking all involve simultaneous interaction. However, with the advent of symbolic communication we learned to utilize longer-term interactions. A physically realized icon, such as a picture, a ring or body decoration, is more permanent. The ultimate extension of this is written language. The coming of telephones, radios and the Internet have obviously extended the interaction distances considerably.

While symbolic communication requires considerable sophistication of interpretation, humans also use signals that can be interpreted more simply. For example, laughter has the same meaning to all humans, but not to other animals. We can make the necessary connection with the emotional state since we can hear and observe others and ourselves laughing – we share the same innate involuntary laugh behavior.

The developmental process that provides the shared groundings for human symbolic communication – cultural language learning – can be seen as an extension of the processes present in our non-human primate ancestors (Hendriks-Jansen, 1996). The major differences being in the complexity due to the sheer number of groundings we need to learn and the intrinsic power of symbolic representation over exclusively indexical and iconic representation. Symbolic representations derive their power because they provide a degree of independence from the symbolic, indexical and iconic references that generated the relationship represented. New symbols can be learnt using language metaphor (Lakoff and Johnson, 1980; Johnson, 1991).

2.4 Symbolic communication and its prerequisites

Evolution does not have the luxury of being able to make simultaneous independent changes to the design of an organism and also ensure their mutual consistency (in terms of the viability of the organism). For this reason, once a particular mechanism has been evolved, it is built upon rather than significantly re-designed to effect a new mechanism. It is only when selection pressures change enough to render things a liability that they may be discarded. This is why layering is observed in natural systems (e.g. Mallot, 1995).

In the examples above, we can perceive a layering of communication mechanisms that are built up as we look at each in turn – from bacteria to humans. Each leveraging the mechanism developed in the previous layer. The ant's use of chemical pheromone trails to implement longer duration interactions is supported by direct-contact chemical communication, pioneered by their distant bacterial ancestors. Wolves also employ this type of communication, which provides an environment that supports the developmental process for learning other shared groundings. The sophistication of such developmental processes is greater in non-humans primates and significantly so in humans. However, even for humans, these processes still leverage the simpler processes that provide the scaffolding of shared iconic and indexical groundings (see Thelen and Smith, 1994; Hendriks-Jansen, 1996).

We believe such layering is integral to the general robustness of biological systems. If a more sophisticated mechanism fails to perform, the lesser ones will still operate. We emulate the layering in the implementation of our system for this reason.

From our examination, the following seem to be necessary for symbolic communication between two agents.

- Some iconic representations in common (e.g. by possessing some physically identical sensory-motor apparatus).
- Either a shared grounding for some indexical representations, a common process that develops shared indexical groundings, or a combination of both (e.g. a mechanism for learning the correlation between icons – such as correlating 'smoke' with 'fire').
- A common process that develops shared symbolic groundings (e.g. mother and infant

'innate' behavior that scaffolds language development – turn-taking, intentional interpretation, mimicking etc.)

Additionally, unless the symbol repertoire is to be fixed with specific processes for acquiring each symbol, it seems necessary to have:

- A mechanism for learning new symbols by communicating known ones (e.g. interpretation and learning through metaphor).

The implementation of this last necessity in a robot system is currently beyond the state-of-the-art. However, the first three are implemented in the cooperative cleaning system, as described in the following section.

3. The System

Our research involved the development of an architecture for behavior-based agents that supports cooperation (Jung, 1998; Jung and Zelinsky, 1999)[2]. To validate the architecture we implemented a cooperative cleaning task using the two Yamabico mobile robots pictured in Figure 2 (Yuta et al., 1991). The task is to clean our laboratory floor space. Our laboratory is a cluttered environment, so the system must be capable of dealing with movable obstacles, people and other hazards.

3.1 The Robots

As we are interested in heterogeneous cooperation, we built each robot with a different set of sensors and actuators, and devised the cleaning task such that it cannot be accomplished by either robot alone. One of the robots, '*Joh*', has a vacuum cleaner that can be turned on and off via software. Joh's task is to vacuum piles of 'litter' from the laboratory floor. As our aim was not to design a high performance cleaning system per se, chopped Styrofoam serves as 'litter'. Joh cannot vacuum close to walls or furniture, as the vacuum is mounted between the drive wheels. It has the capability to 'see' piles of litter using a CCD camera and a video transmitter that sends video to a *Fujitsu MEP tracking vision system*. The vision system uses template correlation, and can match about 100 templates at frame rate. The vision system can communicate with the robot, via a UNIX™ host, over a radio modem. Visual obstacle-avoidance behavior has been demonstrated at speeds of up to 600mm/sec (Cheng and Zelinsky, 1996).

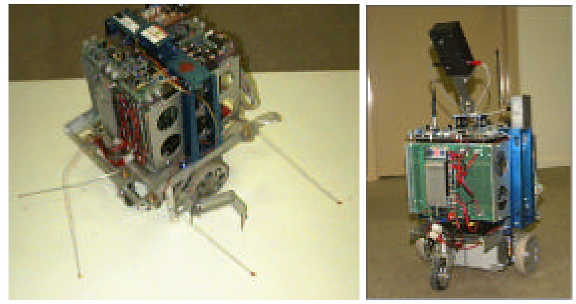

Figure 2 – The two Yamabicos 'Flo' and 'Joh'

The other robot, '*Flo*', has a brush tool that is dragged over the floor to sweep distributed litter into larger piles for Joh to pick-up. It navigates around the perimeter of the laboratory where Joh cannot vacuum and deposits the litter in open floor space. Sensing is primarily using four specifically developed passive tactile 'whiskers' (Jung and Zelinsky, 1996a). The whiskers provide values proportional to their angle of deflection. Both robots are also fitted with ultrasonic range sensors and wheel encoders.

3.2 A layered solution

We implemented the cleaning task by layering solutions involving more complex behavior over simpler solutions. This provides a robust final solution, reduces the complexity of implementation and allows us to compare the system performance at intermediate stages of development.

The first layer involves all the basic behavior required to clean the floor, but does not include any capacity to purposefully navigate, explicitly communicate or cooperate. Flo sweeps up litter and periodically deposits it into piles where it is accessible by Joh. Joh uses the vision to detect the piles and vacuum them up. Therefore, the signaling – depositing litter piles – is implicit in this case, as it is normal cleaning behavior. The interaction is not simultaneous, as Joh doesn't necessarily see the piles as soon as they are deposited. The interaction distance ranges over the size of the laboratory. Flo doesn't interpret the piles of litter as a signal at all – and in fact has no way of sensing them. Joh has a simple interpretation – the visual iconic representation of the pile acts as a releaser to vacuum over it.

[2] See also http://pobox.com/~david.jung/thesis.html

Layer 4
symbolic communication
of litter locations

Layer 3
explicit communication
of litter relative positions

Layer 2
implicit visual communication
of likely litter position

Layer 1
no awarness of each other

Figure 3 - Layered solution

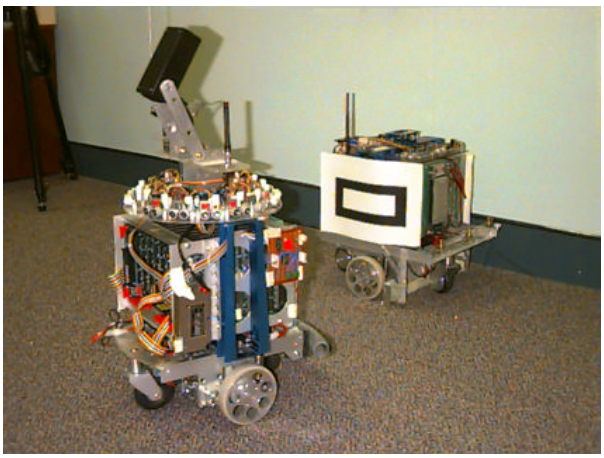

Figure 4 – Joh visually tracking Flo *(no vacuum attached)*

The second layer gives Joh an awareness of Flo. We added the capability for Joh to visually detect and track the motion of Flo. This is another communication mechanism that provides state information about Flo to Joh. In this case, the signaling is again implicit, the interaction distance is visual range and the interaction is simultaneous. Joh uses the visual iconic representation of Flo to ground an indexical reference for the likely location of the pile of litter deposited. Figure 4 shows Joh visually observing Flo via a distinctive pattern. Details of the implementation the visual behavior we employed can be found in (Jung et al., 1998a). A top view of typical trajectories of the robot is shown in Figure 5.

The third layer introduces explicit communication. Specifically, upon depositing a pile of litter, Flo signals via radio the position (distance and orientation) of the pile relative to its body and the relative positions of the last few piles deposited. Flo and Joh both have identical wheel encoders, so we are ensured of a shared grounding for the interpretation of the communicated relative distance and orientation to piles. Although odometry has a cumulative error, this can be ignored over such short distances. The catch is that the positions are relative to Flo. Hence, Joh must transform them to egocentric positions based on the observed location of Flo. If Flo is not currently in view, the information is ignored. A typical set of trajectories is shown in Figure 6.

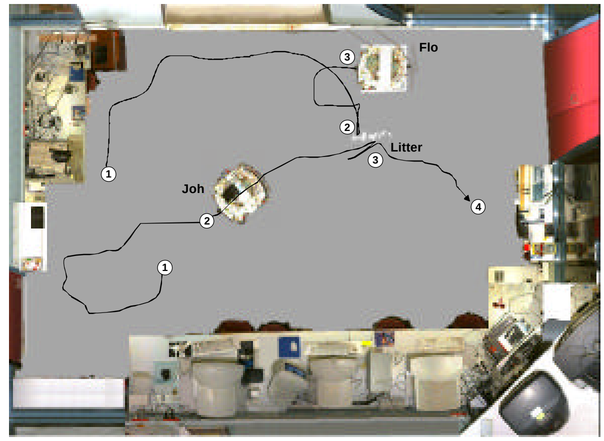

Figure 5 - Typical trajectories when Joh can observe Flo depositing litter (Layer 2)

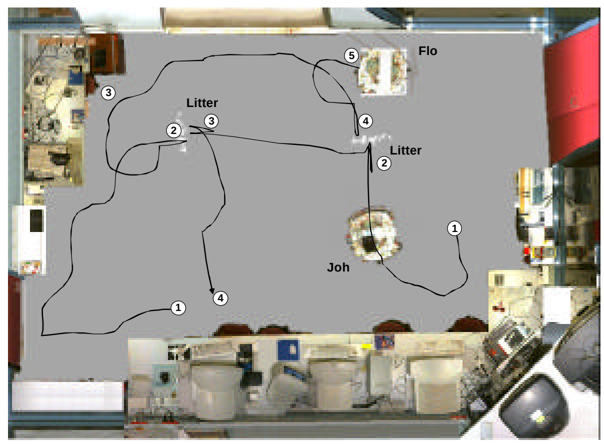

Figure 6 - Typical trajectories when explicit communication is utilized (Layer 3)

The fourth and final layer involves communication of litter locations by Flo to Joh even when Flo cannot be seen. This is accomplished by using a symbolic interpretation for a specific geometric relationship of positions to each other. What is communicated to convey a location is analogous to '*litter position is <specific-geometric-relation-between> <position-A> <position-B> <direction> <distance>*'. The positions are indexical references that are themselves grounded through a shared process, a location-labeling behavior, described below. The distance and direction are in fact raw encoder data, hence an iconic reference, relying on the shared wheel encoders. There is no signal communicated for the symbolic relation itself (like a word), since there is only one symbol in the system, it is unambiguous. Obviously, if more symbols were known, or a mechanism for leaning new symbols available, labels for the symbols would need to be generated and signaled (and perhaps syntax established).

First, we describe the action selection scheme employed, as it is the basis for the navigation and map building mechanism, which in turn is the basis for the location-labeling behavior.

3.3 Action Selection

We needed to design an action selection mechanism that is distributed, grounded in the environment, and employs a uniform action selection mechanism over all behavior components. Because the design was undertaken in the context of cooperative cleaning, we also required the mechanism to be capable of cooperative behavior and communication, in addition to navigation. Each of these requires some ability to plan. This implies that the selection of which action to perform next must be made in the context of which actions may follow – that is, within the context of an ongoing plan. In order to be reactive, flexible and opportunistic, however, a plan cannot be a rigid sequence of pre-defined actions to be carried out. Instead, a plan must include alternatives, have flexible sub-plans and each action must be contingent on a number of factors. Each action in a planned sequence must be contingent on internal and external circumstances including the anticipated effects of the successful completion of previous actions. Other important properties are that the agent should not stop behaving while planning occurs and should learn from experience.

There were no action selection mechanisms in the literature capable of fulfilling all our requirements. As our research is more concerned with cooperation than action selection per se, we adopted Maes' spreading activation algorithm and modified it to suit our needs. Her theory "*models action selection as an emergent property of an activation/inhibition dynamics among the actions the agent can select and between the actions and the environment*" (Maes, 1990a).

3.3.1 Components and Interconnections

The behavior of a system is expressed as a network that consists of two types of nodes – *Competence Modules* and *Feature Detectors*. Competence modules (**CM**s) are the smallest units of behavior selectable, and feature detectors (**FD**s) deliver information about the external or internal environment. A CM implements a component behavior that links sensors with actuators in some arbitrarily complex way. Only one CM can be executing at any given time – a winner-take-all scheme. A CM is not limited to information supplied by FDs – the FDs are only separate entities in the architecture to make explicit the information involved in the action selection calculation.

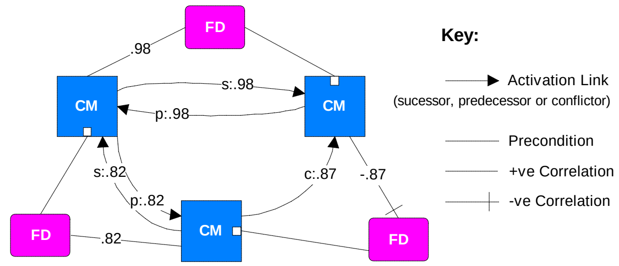

Figure 7 - Network components and interconnections

The graphical notation is shown above where rectangles represent CMs and rounded rectangles represent FDs. Although there can be much exchange of information between CMs and FDs the interconnections shown in this notation only represent the logical organization of the network for the purpose of action selection.

Each FD provides a single *Condition* with a confidence [0…1] that is continuously updated from the environment (sensors or internal states). Each CM has an associated *Activation* and the CM selected for execution has the highest activation from all *Ready* CMs whose activations are over the current global threshold. A CM is *Ready* if all of its *preconditions* are satisfied. The activations are continuously updated by a *spreading activation algorithm*.

The system behavior is designed by creating CMs and FDs and connecting them with *precondition links*. These are shown in the diagram above as solid lines from a FD to a CM ending with a white square. It is possible to have negative preconditions, which must be false before the CM can be *Ready*. There also exist *correlation links*, dotted lines in the figure, from a CM to a FD. The correlations can take the values [-1…1] and are updated at run-time according to a learning algorithm. A positive correlation implies the execution of the CM causes, somehow, a change in the environment that makes the FD condition true. A negative correlation implies the condition becomes false. The designer usually initializes some *correlation links* to bootstrap learning.

Together these two types of links, the precondition links and the correlation links, completely determine how activation spreads thought the network. The other *activation links* that are shown in Figure 7 are determined by these two and exist to better describe and understand the network and the activation spreading patterns. The activation links dictate how activation spreads and are determined as follows.

- There exists a *successor link* from CM *p* to CM *s* for every FD condition in *s*'s preconditions list that is positively correlated with the activity of *p*.
- There exists a *predecessor link* in the opposite direction of every successor link.

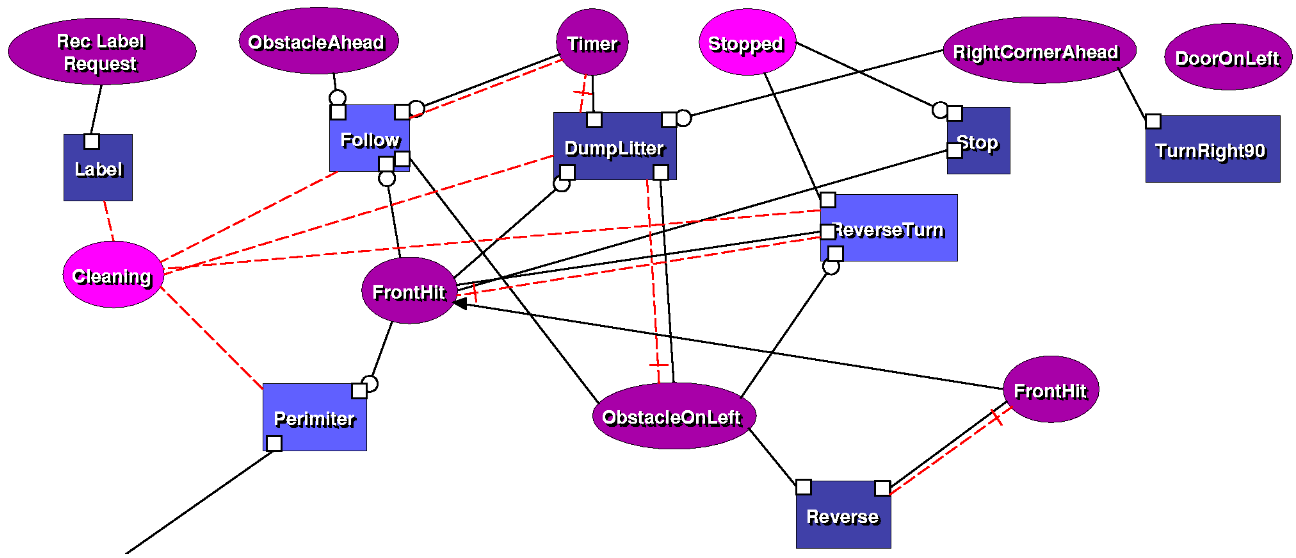

Figure 8 – Partial network for Flo (produced by our GUI).

- There exists a *conflictor link* from CM *x* to CM *y* for every FD condition in *y*'s preconditions list that is negatively correlated with the activity of *x*.

The successor, predecessor and conflictor links resulting from the preconditions and correlations are shown in Figure 7.

In summary, a CM *s* has a predecessor CM *p*, if *p*'s execution is likely to make one of *s*'s preconditions true. A CM *x* has a conflictor CM *y*, if *y*'s execution is likely to make one of *x*'s preconditions false.

3.3.2 The Spreading of Activation

A rigorous description of the spreading activation algorithm is beyond the scope of this paper. The algorithm has been detailed in previous publications (Jung, 1998; Jung and Zelinsky, 1999). The activation rules can be more concisely described in terms of the activation links. The main spreading activation rules can be simply stated:

- *Unready* CMs increase the activation of predecessors and decrease the activation of conflictors, and
- *Ready* CMs increase the activation of successors.

In addition, these special rules change the activation of the network from outside in response to goals and the current situation:

- Goals increase the activation of CMs that can satisfy them and decrease the activation of those that conflict with them, and
- FDs increase the activation of CMs for which they satisfy a precondition.

To get a feel for how it works, we describe part of a network that implements the cleaning task for Flo, as shown in Figure 8. With some of the components shown, a crude perimeter-following behavior is possible.

The rectangles are basic behaviors (CMs), the ovals feature detectors (FDs), and only the correlation and precondition links are shown (the small circles indicate negation of a precondition). The goal is `Cleaning`. This occurs when Flo roughly follows the perimeter of the room by using `Follow` to follow walls and `ReverseTurn` to reverse and turn away from the perimeter when an obstacle obstructs the path. Periodically the litter that has accumulated in the sweeper is deposited away from the perimeter by `DumpLitter`.

The spreading activation algorithm 'injects' activation into the network CMs via goals and via FDs that meet a precondition. Therefore, the `Cleaning` goal causes an increase in the activation of `Follow`, `DumpLitter` and `ReverseTurn`. Suppose Flo is in a situation where its left whiskers are against a wall (`ObstacleOnLeft` is true) and there are no obstacles in front (`ObstacleAhead` and `FrontHit` both false). In this case, the activation of `Follow` will be increased by all the FDs in its precondition set (including `Timer` which is false before being triggered). Being the only CM *ready*, it is scheduled for execution until the situation changes. Once the `Timer` FD becomes true, `Follow` is no longer ready, but `DumpLitter` becomes ready and is executed. `Follow` and `DumpLitter` also decrease each other's activation as they conflict – each is correlated with the opposite state of `Timer`.

Although, the selection of CMs in this example depends mainly on the FD states, when the selection of CMs depends more on the activation spread from other CMs, the networks can exhibit 'planning' – as Maes has shown. This is the basis for action planning in our networks, and gives rise to path planning as will be described below.

From the rules we can imagine activation spreading backward through a network, from the goals, through CMs with unsatisfied preconditions via the precondition links until a *ready* CM is encountered. Activation will tend to accumulate at the ready CM, as it is feeding activation forward while its successor is feeding it backward. Eventually it may be selected for execution, after which its activation is reset to zero. If its execution was successful, the precondition of its successor will have been satisfied and the successor may be executed (if it has no further unsatisfied preconditions). We can imagine multiple routes through the network, activation building up faster via shorter paths. These paths of higher activation represent 'plans' within the network. The goals act like a 'homing signal' filtering out through the network and arriving at the current 'situation'.

One important difference between our and Maes' networks is that in ours the flow of activation is weighted according to the correlations – which are updated continuously at run-time according to previous experience. The mechanism for adjusting the correlation between a given CM-FD pair is simple. Each time the CM becomes active, the value of the FD's condition is recorded. When the CM is subsequently deactivated, the current value of the condition is compared with the recorded value. It is classified as one of: *Became True*, *Became False*, *Remained True* or *Remained False*. A count of these cases is maintained (B_t, B_f, R_t, R_f). The correlation is then:

$$corr = \frac{(2B_t + R_t)}{2N} - \frac{(2B_f + R_f)}{2N}$$

Where the total samples $N = B_t + B_f + R_t + R_f$

To keep the network plastic, the counts are decayed so recent samples have a greater effect than historic ones.

3.4 Navigation and map building

3.4.1 Spatial and Topological path planning

There are two main approaches to navigational path planning. One method utilizes a geometric representation of the robot environment, perhaps implemented using a tree structure. Usually a classical path planner is used to find shortest routes through the environment. The distance transform method falls into this category (Zelinsky et al., 1993). These geometric modeling approaches do not fit with the behavior-based philosophy of only using categorizations of the robot-environment system that are natural for its description, rather than anthropocentric ones. Hence, numerous behavior-based systems use a topological representation of the environment in terms only of the robot's behavior and sensing (e.g. see Mataric, 1992). While these approaches are more robust than the geometric modeling approach, they suffer from non-optimal performance for shortest path planning. This is because the robot has no concept of space directly, and often has to discover the adjacency of locations.

Consider the example below, where the robot in (a) has a geometric map and its planner can directly calculate the path of least Cartesian distance, directly from A to D. However, the robot in (b) has a topological map with nodes representing the points A, B, C and D, connected by a *follow-wall* behavior. Since it has never previously traversed directly from A to D, the least path through its map is A-B-C-D.

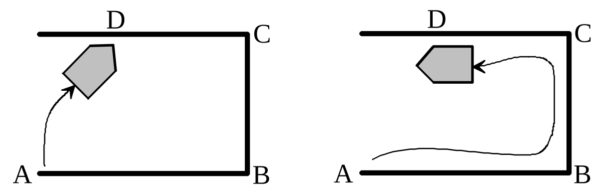

Figure 9 – (a) Geometric vs (b) Topological Path Planning

Consequently, our aim was to combine the benefits of geometric and topological map representations in a behavior-based system using our architecture.

3.4.2 A self-organizing map

In keeping with the behavior-based philosophy, we found no need to explicitly specify a representation for a map or a specific mechanism for path planning. Instead, by introducing the key notion of *location feature detectors* (location FDs), the correlation learning and action selection naturally gave rise to map building and path planning – for 'free'.

A location feature detector is a component of our architecture specialized to respond when the robot is in a particular location (the detector's characteristic location). We employ many detectors and the locations to which they respond are non-uniformly distributed over the laboratory floor space. Each location FD contains a vector $\boldsymbol{v}$, whose components are elements of the robot state vector:

$\boldsymbol{x} = [g,s,fds]$

where $\boldsymbol{g} = (x,y,q)$ global odometry coordinates

$\boldsymbol{s}$ = sensor values

$\boldsymbol{fds}$ = non-location FD values

The variable $\boldsymbol{g}$ contains global Cartesian coordinates and orientation estimated from wheel encoders and a model of the locomotion controller. The sensors include ultrasonic range readings and in Flo's case, tactile whisker values. The ***fds*** component contains the condition values of all FDs in the system, except for the location FDs themselves. For example, in Joh's case this includes visual landmark FDs.

The condition confidence value of each location FD is updated by comparing it to the current state of the robot's sensors and other non-location FDs. A weighted Euclidean norm N_w is used – with the (x,y) coordinate weights dominating.

$$c_i = 1 - N_w(\boldsymbol{v}, \boldsymbol{x})$$

Hence, the vector of the location FD whose condition is true with highest confidence is considered to represent the 'current location' of the robot. The detectors are iconic representations of locations (see Figure 10).

The location FD vectors $\boldsymbol{v}$ are initialized such that the (x,y) components are distributed as a regular grid over the laboratory floor space, and the other components are randomly distributed over the vector space. During operation of the system, the location FD vectors are updated using Kohonen's *self-organizing map* (SOM) algorithm (Kohonen, 1990). This causes the spatial distribution of the location FD vectors to approximate the frequency distribution of the robot's state vector over time. Figure 16 shows how the detectors have organized themselves to represent one of our laboratories. One useful property of a SOM is that it preserves topology – nodes that are adjacent in the representation are neighboring locations in the vector space.

Since the location FD vectors $\boldsymbol{v}$ are continuously matched with the robot state vector $\boldsymbol{x}$, in which the (x,y) coordinates are estimated via odometry, there is a major drawback. The odometry error in (x,y) is cumulative. We remedy this by updating the robot state vector coordinates. Specifically, the system has feature detectors for various landmark types that are automatically correlated with the location FDs by the correlation learning described above. If it should happen that a landmark FD becomes true with high confidence that is strongly correlated with a location FD *neighboring* the location FD for the 'current location', then the state vector (x,y) component is updated. The coordinates are simply moved closer to the coordinates of the location FD to which the landmark is correlated. Assuming the landmarks don't move over moderate periods, this serves to keep the location FD (x,y) components registered with the physical floor space.

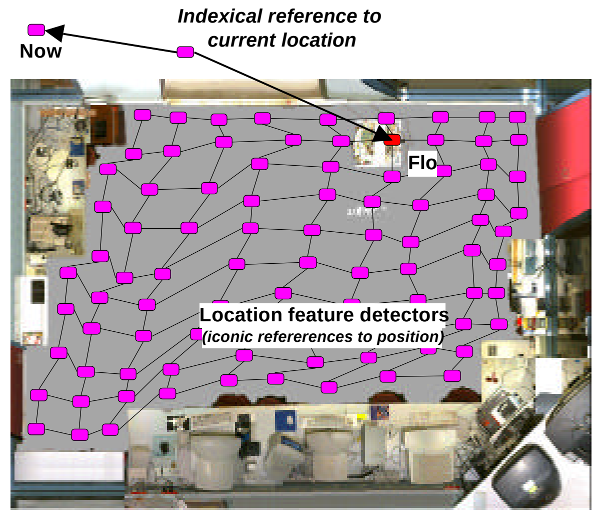

Figure 10 – Schematic of location detector SOM and current location index

The system also maintains an indexical reference that represents the robot's current location. Recall that an indexical reference is a correlation between icons. The robots each have a sense of time – in terms of the ordering relation between sensed events (which is shared to the extent that the ordering of external events is perceived to be the same by both robots). Hence, the current location index is an association between the most active location detector and the current time.

It is clear this mechanism fulfills our requirement for spatial mapping. The topological mapping derives again from the correlation learning in the architecture. Specifically, the system learns by experience that a particular behavior can take the robot from one state to another – for example by changing the current location index in a consistent way. Over time, behavior such as follow-wall becomes correlated with the start and end locations of a wall segment. The spreading activation will cause the behavior to be activated when the system needs to 'plan' a sub-path from the start to the end. Similarly, simple motion behavior becomes correlated with moving the robot from one location to one of its neighbors.

3.4.3 Navigation

Once we have feature detectors that respond to specific locations, it is straightforward to add spatial and topological navigation. Each time a behavior (a CM) is activated, the identity of the current location FD before and after its execution is recorded. A new instance of the CM is created, and initialized with the 'source' location FD as a precondition and the 'destination' as a positive correlate. Hence, the system remembers which behavior can take it from one specific location to another. If the CM does not consistently do this, its correlation with the destination location FD will soon fall. If it falls to zero, the CM is removed from the network. Changes in the environment also cause correlations to change, thus allowing the system to adapt.

With this mechanism, the system learns topological adjacency of locations in terms of behavior. For example, if the activation of the `Follow` CM consistently takes the robot from the location FD corresponding to the start of a wall, to the end of the wall, then the links shown below will be created.

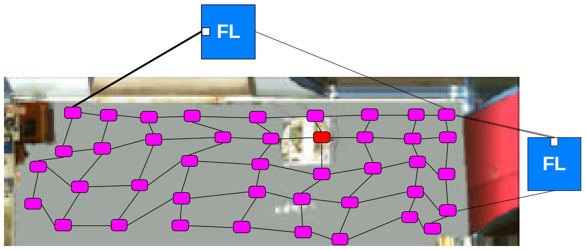

Figure 11 - Behavioral adjacency of locations via Follow (FL)

The spreading activation algorithm for action selection is able to plan a sequence of CM activations to achieve navigation between any arbitrary locations.

Spatial navigation is achieved by initializing the network so that a simple `Forward` behavior links each location FD with its eight neighbors in both directions. Hence, initially the system 'thinks' it can move in a straight line between any locations that are neighbors in the SOM. If presence of an obstacle blocks the straight-line path from one location to its neighbor, then this will be learnt through a loss of correlation between the corresponding `Forward` CM and the 'destination' FD. The mechanisms described here for map building and navigation are presented in detail in (Jung, 1998; Jung and Zelinsky, 1999a).

3.5 A shared grounding for locations

For layer 4 of the implementation, we wanted to add the capability for Flo to communicate the locations of litter piles in a more general way. In such a way that it would be useful to Joh if Flo were not in view or even in another room. In the system as described thus far, Flo and Joh do not share any representations except the iconic representations of their shared sensors (odometry and ultrasonic). The location feature detectors may be correlated with visual landmarks in Joh's map, and whisker landmarks in Flo's (among other information).

Hence, before we can communicate Flo's representation for location we need a procedure to establish a shared grounding with Joh. For this purpose, we have implemented a *location labeling* procedure. Location labeling is essentially behavior whereby Flo teaches Joh a location by empirical demonstration. It proceeds as follows.

If Joh is tracking Flo in its visual field at a particular time and there are no previously labeled locations near by, then Joh signals Flo indicating that Flo's current location should be labeled. Although an arbitrary signal could be generated and communicated to serve as a common labeling icon for the location, in this specific case no signal is necessary. Because there are only two robots, the time ordering of the labeling procedures is identical to each. Hence, a time ordered sequence number maintained by each serves as the labeling icon with a shared grounding. The first location is labeled '1st Label', the next '2nd Label', etc. If Joh receives a confirmation signal from Flo, it associates the label icon with Flo's current location. Joh calculates Flo's location based on its own location and a calculation of Flo's range from visual tracking information. Flo also labels its own location index in the same way. This procedure creates an indexical representation of specific locations that are associations between a location detector icon and the label icon (the shared sequence number). Although the locations themselves are not represented using the same icons by both Flo and Joh, they represent the same physical location. Figure 12 shows the situation after the labeling procedure has occurred four times (the symbol is explained below).

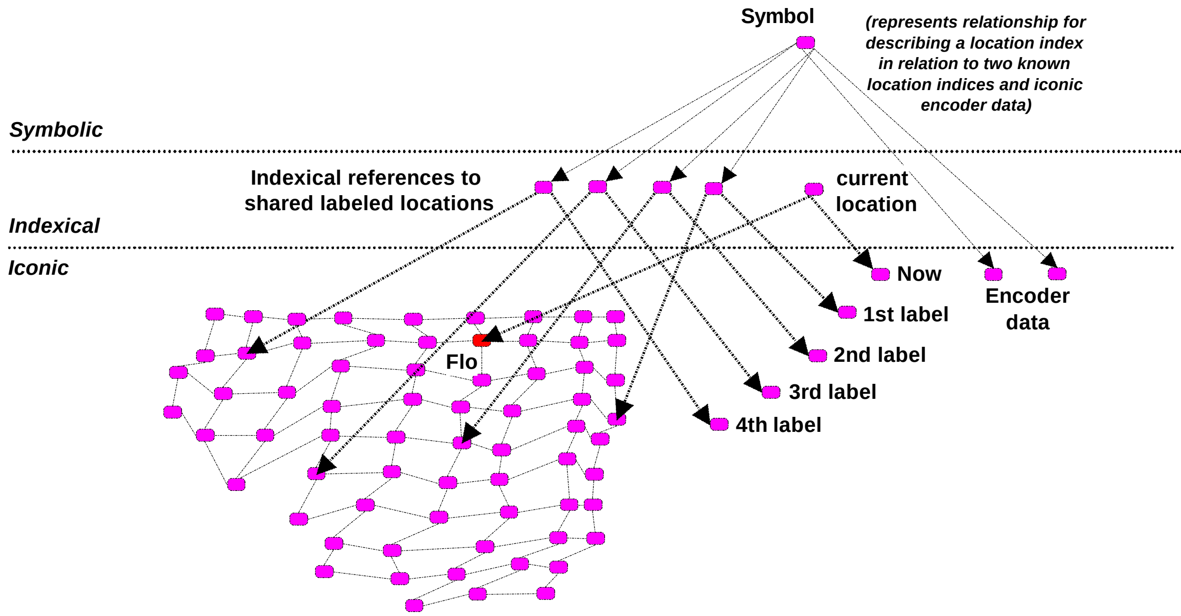

Figure 12 – Shows Iconic references that represent sensory data, Indexical references that associate pairs of icons (a label with a location) and a symbol (see text). The fine lines between location feature detectors show their adjacency in the SOM; the pairs of arrow headed lines from indexical references define which two icons they associate; and the two sets of arrow headed lines from the symbol designate two 'exemplars' (see text).

3.6 A symbol for a relationship between locations

The next step is to endow both Joh and Flo with the ability to represent an arbitrary location in relationship to already known locations. Recall that a symbol is defined as a relationship between other symbolic, indexical and iconic references.

Ideally, symbols should be learnt, as in biological systems. The relationship a symbol represents is a generalization from a set of observed 'exemplars' – specific relationships between other symbols, indices and icons. How this can be accomplished is still an open research area. For this reason, and because we only need a single symbol that will not be referenced by higher-level symbols, we chose to simply provide the necessary relationship. We can consider the symbol a 'first-level symbol', as it is not dependent on any other symbols, but grounded directly to iconic and indexical representations. As symbol systems go, ours is as impoverished as it can be.

The relationship represented by the symbol is between two known location indices and a distance and orientation in terms of wheel encoder data. The two known locations define a line segment that provides an origin for position and orientation. The wheel encoder data then provides a distance and orientation relative to this – which together defines a unique location (see Figure 13). For example, a pile could be specified as being approximately 5m away from the 2^{nd} labeled location at an angle of 30^{o} relative to the direction of the 1^{st} labeled location from the 2^{nd}. The top of Figure 12 shows the symbol in the context of the overall system.

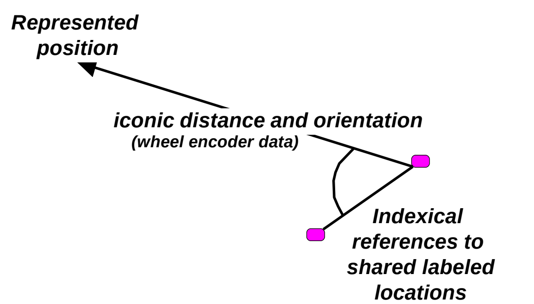

Figure 13 - Schematic of the *<specific-geometric-relation-between>* symbol used to communication locations

3.7 Symbolic communication

Finally, we are in a position to see how a location can be symbolically communicated from Flo to Joh. With a particular pile location in mind, Flo first calculates the representation for it using the symbolic relationship above. It selects the two closest locations, previously labeled, as the indexical references and computes the corresponding iconic wheel encoder data that will yield the desired pile location. This information is then signaled to Joh by signaling the labels for each of the known locations in turn, followed by the raw encoder data. This signal is grounded in both robots, as the labels were grounded through the location labeling procedure, and the wheel encoders are a shared sense. Hence, the meaning is preserved. Joh can recover the location by re-grounding the labels and reversing the computation.

3.8 Results

The typical trajectories in Figure 14 show that Joh is able to successfully vacuum the litter in the pile to the left. This occurs after the location of the pile has been communicated symbolically by Flo. The pile was initially obscured by the cardboard box, but Joh was able to correctly compute its location and plan a path around the box using its map. This can be contrasted with the layer 3 solution shown in Figure 6, where no symbolic communication or map was utilized. If the box were blocking the straight-line path to the litter pile in that case, Joh would not have been able to navigate to within visual range to locate it.

As the system was not designed as a floor cleaning system per-se, rigorous experiments to record its cleaning performance were not conducted. However, we did run experiments that seem to show that the addition of symbolic communication does improve cleaning performance. We expect this intuitively, as the governing factor in vacuuming performance is the path length between litter piles. The ability to navigate purposively from one known litter pile location to the next, instead of having to rely on an obstacle free path, or chance discovery of the pile locations, shortens the average path length.

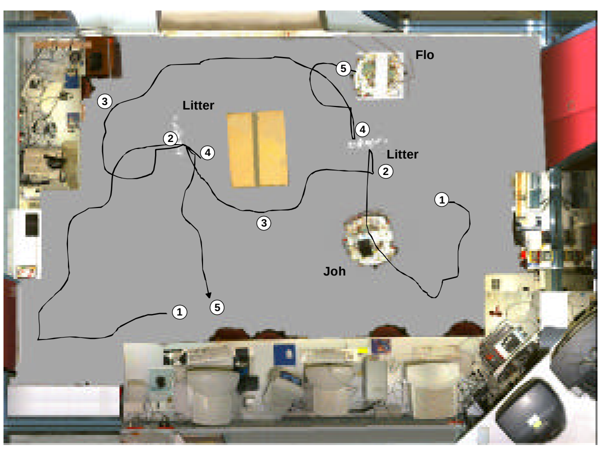

Figure 14 - Typical trajectories during cooperation

We also ran experiments utilizing each layer in turn (including the lower ones on which it builds). We recorded the percentage of the floor cleaned every two minutes from 3-15 minutes. It was difficult to run all of the experiments consistently for more than 15 minutes due to problems with hardware reliability. The results are plotted in Figure 15. Initially, about 30% of the 'litter' was distributed around the perimeter and the remainder scattered approximately uniformly over the rest of the floor. The percentage cleaned was estimated by dividing the floor into a grid and counting how many tiles had been cleaned.

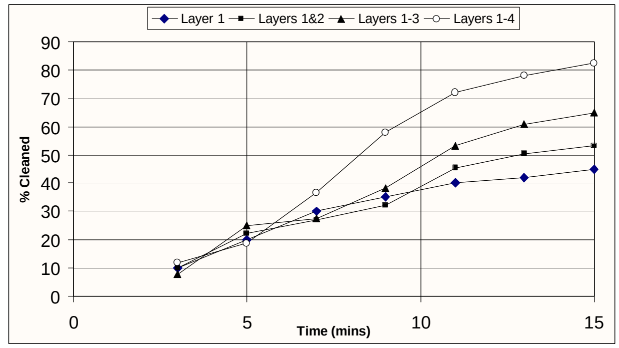

Figure 15 - Performance of layered cleaning solutions

Clearly, the addition of each layer improves the cleaning performance. In particular, layer 4, utilizing symbolic communication, initially falls behind as some time is used to perform location labeling rather than cleaning. This starts to pay off later after a number of locations have been labeled.

This experiment also shows the robustness gained by layering the solution. The implementation of layer 1 is robust due to its simplicity. If any of the mechanisms employed in the subsequent layers were to fail, we have demonstrated that the system will continue to perform the cleaning task, although not as quickly.

4. A Critical Examination

4.1 The limitations of our system

The are two obvious limitations to the approach we have described for developing grounded symbolic communication between robots. The first is that the common process by which a shared symbol grounding is developed is the design process. That is, the shared grounding was established by identical design and implementation of the mechanism for its interpretation. This is an impractical way to develop sophisticated symbol systems, as the mechanism for the interpretation of each symbol must be designed in turn.

Is this just a practicality problem, or it is impossible in principle? When we designed the system, we believed that it was possible, if impractical, to build general symbol systems in this way – by explicitly designing the process of interpretation for each symbol. We hypothesized that all that was missing was a mechanism to learn the symbolic representations – to effectively automate the process. However, we argue below that is it in fact impossible in principle (for all but the simplest symbol systems – like the one presented).

The second obvious limitation is a related one. The approach doesn't include a mechanism for learning new symbols, even if it had an existing symbol repertoire designed in.

4.2 Symbols revisited

Our definition of *grounded* from section 2.2 contained a hidden assumption. We defined a symbol to be grounded if its interpretation required following references that all eventually lead to icons. Recall that, symbols and the structure of their relationships to each other and to indices and icons, is a linguistic one. It is the empirically observed structure of the *signals* that humans generate and interpret. This grammatical structure of spoken and written language is a relatively persistent one (ignoring the fact that languages change slowly over time). The hidden assumption, which we now believe to be incorrect, was that this somehow implies that a similarly persistent analogous structure must be present within the mind of the humans that generate signals conforming to the structure. That is, just because there is a relatively persistent symbolic system present in human cultural artifacts – such as books, paintings, buildings, music, etc. – this does not imply that any symbol system persists within the human mind. It was with this invalid assumption that we proceeded to construct just such a symbol system within the robots, by representing and relating icons, indices and symbols.

We believe there is ample evidence that no persistent symbolic structure within the human mind that mirrors the structure of human language exists – but this remains to be seen. Dennett has argued strongly against the idea of a Cartesian theater – a place in the mind where all the distributed information is integrated for a central decision-maker (Dennett, 1993). It seems that distributed information about the external world (possibly contradictory) need not be integrated unless a particular discrimination is necessary for performance (for example to speak or behave). Even then, only the information necessary for the discrimination need be integrated.

Even if humans don't use the equivalent of a persistent cognitive grammar to reason about the world, why can't robots use one?

4.3 Symbolic representation is not situated

A symbol represents a discrete category in the continuous space of sensory-motor experience. Hence it defines a boundary such that points in the space lie either within the category or outside of it – there are no gray areas. Therefore, a symbol system is a way of characterizing sensory-motor experience in terms of membership of the categories it defines. Symbols derive their power by conferring a degree of independence from the context dependent, dynamic and situated experiences from which they are learnt. This allows symbolic communication to preserve its meaning when the interaction is extended in time (e.g. the period between these words being written and you reading them).

Suppose we build a robot for a particular task that necessitates symbolic communication, and endow it with a symbolic representation system according to the approach we have outlined, whereby static symbol groundings are designed in. The robot is *situated* in the sense that the task for which is it designed provides a context for its interaction with the environment (from the theory of situated action – Mills, 1940; Suchman, 1987). The robot is an embodied agent and has a grounded symbol system. It satisfies the criteria of the *physical grounding hypothesis* (Brooks, 1990).

We argue that this approach to building a robot will not necessarily work, except in the simplest cases. The task in which the robot is situated dictates the discriminations it must make in order to behave appropriately – it must behave in terms of its *affordances* (Gibson, 1986). Since the discriminations it can make are determined by the categories defined by its symbol system, which is necessarily static, it will only work if the task very specific – ensuring the appropriate discriminations don't change. This is precisely the situation in which our system operates – in the situated context defined by a statically specified cleaning task.

A robot capable of operating flexibly in a dynamic situated context must continually adapt the discriminations it makes. If using a symbolic representation system, this implies the categories defined by the symbols, and hence the meaning of the symbols themselves, must change[3]. However, a dynamic symbol system looses its power for communication – one of the main reasons for endowing the robot with a symbol system in the first place.

Consequently, a robot that utilizes a static symbolic representation system (like the one we presented) cannot be *situated* if its task is to behave flexibly in a dynamic context. Hence, our approach of designing in the robot's symbol groundings does not scale from systems designed to achieve simple specific tasks, to more general flexible behavior.

We also see a more pragmatic way in which larger symbol systems built via our approach can become unsituated. In order to manage complexity in the design process, we often structure a system by categorizing and apply linguistic labels to design components (i.e. we need to name elements of our designs). Although this activity is logically independent from the way the system

[3] It may be possible in principle for an agent to use a static symbol system that covers all possible categorizations and hence accommodates any possible discrimination needed for appropriate behavior in any situated context. However, we dismiss this as impossible in practice due to computation intractability.

functions, the anthropocentric groundings we use in our interpretation of the linguistic labels inevitably effect the design.

For example, by naming a behavior component `WallFollowing`, we may accidentally allow hidden assumptions from our understanding of 'walls' to come into play, despite being aware of this pitfall. If the robot possesses anything that could be called a concept for a 'wall', it is surely impoverished compared to our human understanding of 'walls'. We contend that avoiding this pitfall becomes harder, to the point of practical impossibility, as the symbol systems become more complex and the discrepancy between our labels and the robot's representations grow.

4.4 Adaptive Symbol Grounding Hypothesis

There is increasing evidence that humans do not reason about the world and behave using symbolic representations (Hendriks-Jansen, 1996 provides a thorough argument). Instead, like other biological systems, we represent[4] the world in terms of changing affordances – dictated by our situatedness. We make only the discriminations necessary to behave appropriately. The symbols we use to communicate seem to be generated during language production and interpretation by a dynamic process that grounds them in our adaptive internal representations while preserving their static, public, statistically persistent meaning. Hence, the symbols we generate are influenced by our situated representations during production and they have the power to influence them during interpretation. The symbolic representations themselves are only transient. We refer to this conception as the *Adaptive Symbol Grounding Hypothesis*.

By this conception, we envisage the process of learning new concepts as follows. A process within the emitter wishing to communicate a new concept dynamically generates a transient symbolic representation that best approximates it by matching the internal representation with learnt static linguistic symbol relationships. This structure is reflected in the signal. The interpretation process within the receiver causes a similar transient symbolic structure to emerge. Again, an approximate match is made between the symbolic structure and the internal representation – which influences the representations. In this case, the influence causes a new concept to be discovered. The symbolic structure provides the scaffolding necessary to get the receiver thinking in the right way to discover the new concept.

[4] We do not mean to imply that biological agents represent the world to themselves. Of course any observations of the internal states of an agent can be said to represent something – if we as scientific observers interpret it, it represents something to us.

So the essential points of the *Adaptive Symbol Grounding Hypothesis* can be summarized as follows.

- The persistent relationships between icons, indices and symbols that comprise the hierarchical structure of language (e.g. grammar) are only observed in the communicated signals.
- Agents engaging in symbolic communication do not need to maintain an explicit representation analogous to the symbolic structure of the language.
- Symbol grounding is transient and adaptive. Explicit symbolic representations and their situated groundings only persist during the generation and interpretation of the signals of symbolic communication. The specific groundings with which icons for particular symbols are associated depend upon a history of use. The mapping adapts both to the immediate context and to track long-term common usage within a community of language users.

4.5 Implication for cooperative robotics

In the future, we will require increasingly complex tasks to be carried out by multi-robot teams. Hence, the behavioral sophistication of the individual robots will be greater. If we wish to engineer multi-robot systems that can cooperate in complex ways, they will eventually require symbolic communication.

The *Adaptive Symbol Grounding Gypothesis* implies that all symbols are learnt. Hence, we advocate the ubiquitous use of learning in engineering all robotic systems. Without it, we don't believe symbolic communication of significance is possible.

Multi-robot systems are usually classified as either homogeneous or heterogeneous. This is usually based upon physical attributes, such as sensors and actuators; but can be equally applied to the computational and behavioral ability of the robots. A robot system is classified as heterogeneous if one or more agents are different from the others. Balch proposes a metric to measure the diversity in multi-robot systems he calls *social entropy* – which also recognizes physically identical robots that differ only in their behavioral repertoire (Balch, 1997).

If robots are engineered with an emphasis on learning and are consequently more a product of their experience, as we suggest above, then even physically homogeneous teams will have significant social entropy. The teams will necessarily be heterogeneous in terms of their representation of the world and hence behavior. Therefore, we don't envisage homogeneous multi-robot systems playing a large role in the cooperative robotics domain in the long term.

5. Summary

In the first part of the paper, we defined what we mean by *grounded* and provided a framework for talking about symbols in terms of indexical and iconic references. We also introduced the classification scheme for communication involving the characteristics *interaction distance*, *interaction simultaneity*, *signaling explicitness* and *sophistication of interpretation*. We discussed cooperation and communication in bacteria, ants, wolves, primates and humans in these terms to deduce some prerequisites for symbolic communication.

If we are not interested in preserving the meaning of a signal between emitter and receiver, then the implementation is straightforward. If we wish to preserve meaning, then we have to ensure a shared grounding between the agents. In the case of iconic representations, as they are essentially grounded directly in sensory information, this can only be ensured if the sensors are identical between the agents. In the case of indexical and symbolic representations, a specific mechanism for establishing a shared grounding is needed. For indexical representations, an empirical demonstration can serve to ground them to appropriate icons. The *location labeling* procedure we implemented on our robots takes this form.

We described the implementation of the cooperative cleaning system, including the spreading activation action-selection mechanism and purposive navigation in order to provide an understanding for the communication mechanism. The symbolic communication relies on:

- the shared grounding of icons through common sensors,
- the shared grounding for locations, developed through a specific process – the location labeling behavior, and
- the shared grounding for the symbol representing a specific relationship between locations – provided by design.

In the final part of the paper, we critically examined the system and its limitations. Specifically, one obvious limitation is that the system only contains a single symbol, and it was provided at design time – with no mechanism for learning further symbols. By looking again at the notion of a *symbol* we were able to understand that this approach cannot scale to larger symbol systems.

We argued that situated, embodied agents cannot use symbolic representations of the world to interactively behave in it. The *Adaptive Symbol Grounding Hypothesis* was introduced as an alternative conception for how symbol system might be used in situated agents. Finally, we concluded that symbol grounding must be learnt. Consequently, we advocate the ubiquitous use of learning in heterogeneous multi-robot systems, because without it symbolic communication is not possible. We believe this would be a severe limitation to the sophistication of cooperation in the future.

References

Arkin, R. C. and Hobbs, J. D. 1992b. *Dimensions of Communication and Social Organization in Multi-Agent Robotic Systems,* Proc. Simulation of Adaptive Behavior 92, Honolulu, HI.

Balch, Tucker and Arkin, Ronald C. 1994. *Communication in Reactive Multiagent Robotic Systems,* Autonomous Robots, 1, 27-52, Kluwer Academic Publishers, Boston.

Balch, T. 1997. *Social Entropy: a New Metric for Learning Multi-robot Teams,* Proc. 10th International FLAIRS Conference (FLAIRS-97).

Bond, Alan H. 1996. *An Architectural Model of the Primate Brain,* Dept. of Computer Science, University of California, Los Angeles, CA 90024-1596.

Brooks, Rodney A. 1990. *Elephants Don't Play Chess,* Robotics and Autonomous Systems 6, pp3-15.

Brooks, Rodney A. 1991. *Intelligence Without Reason,* MIT AI Lab Memo 1293. Prepared for Computers and Thought, IJCAI-91.

Bruner, J. S. 1982. *The Organisation of Action and the Nature of Adult-Infant Transaction,* in M. von Cranach and R. Harré (eds.), "The Analysis of Action", pp 313-328, Cambridge: Cambridge University Press.

Cao, Y. Uny, Fukunaga, Alex S., Kahng, Andrew B. and Meng, Frank 1995. *Cooperative Mobile Robotics: Antecedents and Directions,* IEEE 0-8186-7108-4/95.

Cheney, D. L., Seyfarth, R. M. 1990. *How Monkeys see the world,* University of Chicago Press.

Cheng, Gordon and Zelinsky, Alexander 1996. *Real-Time Visual Behaviours for Navigating a mobile Robot,* Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), vol 2. pp973.

Crespi, B. J. and Choe, J. C. (eds.) 1997. *The Evolution of Social Behaviour in Insects and Arachnids,* Cambridge, Cambridge University Press.

Deacon, Terrance 1997. *The Symbolic Species: The co-evolution of language and the human brain,* Penguin Books, ISBN 0-713-99188-7.

Dennett, Daniel C. 1993. *Consciousness Explained,* Penguin Books, ISBN 0-14-01.2867-0.

Dudek, G., Jenkin, M., Milios, E. and Wilkes D. 1993. *A taxonomy for swarm robots,* Proc. International Conference on Intelligent Robots and Systems (IROS), San Francisco, CA, pp1151-1156.

Ford, K. and Hayes, P. (eds) 1991. *Reasoning Agents in a Dynamic World: The Frame Problem,* JAI Press.

Gibson, James J. 1986. *The Ecological Approach to Visual Perception,* Lawrence Erlbaum Associates, London, ISBN 0-89859-958-X.

Hendriks-Jansen, Horst 1996. *Catching Ourselves in the Act,* A Bradford Book, MIT Press, Cambridge, Massachusetts, *ISBN 0-262-08246-2.*

Johnson, M. 1991. *Knowing through the body,* Philosophical Psychology, 4, 3-18.

Jung, David and Zelinsky, Alexander 1996a. *Whisker-Based Mobile Robot Navigation,* Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), vol 2. pp497-504.

Jung, David 1998. *An Architecture for Cooperation among Autonomous Agents,* PhD Thesis, Intelligent Robotics Laboratory, University of Wollongong, Australia. For further information on this research see the web site
http://pobox.com/~david.jung/thesis.html

Jung, David, Heinzmann, Jochen and Zelinsky, Alexander 1998a. *Range and Pose Estimation for Visual Servoing on a Mobile Robotic Target*, Proc. IEEE International Conference on Robotics and Automation (ICRA), vol 2. pp1226-1231, Leuven, Belgium.

Jung, David and Zelinsky, Alexander 1999. *An architecture for distributed cooperative planning in a behaviour-based multi-robot system*, Journal of Robotics and Autonomous Systems (RA&S) 26, pp149-174.

Jung, David and Zelinsky, Alexander 1999a. *Integrating Spatial and Topological Navigation in a Behavior-Based Multi-Robot Application*, proceedings of the International Conference on Intelligent Robots and Systems (IROS99), pp323-328, Kyongju, Korea.

Kaiser, D. and Losick, R. 1993. *How and Why Bacteria Talk to Each Other,* in Cell, Vol. 73, No. 5, pp873-885.

Kohonen, Teuvo 1990. *The self-organising map,* Proceedings of IEEE, 78(9):1464-1479.

Kube, C. Ronald, Zhang, Hong 1994. *Collective Robotics: From Social Insects to Robots*, Adaptive Behaviour, Vol. 2, No. 2, 189-218.

Lakoff, G. and Johnson, M. 1980. *Metaphors we Live By,* Chicago: Chicago University Press.

Maes, P. 1990a. *Situated Agents Can Have Goals.,* Designing Autonomous Agents. Ed: P. Maes. MIT-Bradford Press, 1991. ISBN 0-262-63135-0. Also published as a special issue of the Journal for Robotics and Autonomous Systems,Vol. 6, No 1, North-Holland.

Mallot, Hanspeter A. 1995. *Layered Computation in Neural Networks,* The Handbook of Brain Theory and Neural Networks, Michael A. Arbib (ed.), MIT Press, Bradford Books, ISBN 0-262-01148-4, pp513.

Mataric, Maja J. 1992. *Integration of Representation Into Goal-Driven Behavior-Based Robots,* in IEEE Transactions on Robotics and Automation, Vol. 8, No. 3, 304-312.

Mataric, Maja J. 1992a. *Behavior-Based Systems: Key Properties and Implications,* n Proceedings, IEEE International Conference on Robotics and Automation, Workshop on Architectures for Intelligent Control Systems, Nice, France, 46-54.

Mataric, Maja J 1997a. *Using Communication to Reduce Locality in Distributed Multi-Agent Learning,* Proceedings, AAAI-97, Providence, Rhode Island, July 27-31, pp643-648.

Mills, C. W. 1940. *Situated actions and vocabularies of motive,* American Sociological Review, 5:904-913.

Newell, A. and Simon, H. A. 1972. *Human problem solving,* Englewood Cliffs, NJ: Prentice-Hall, Inc.

Parker, Lynne E. 1995. *The Effect of Action Recognition and Robot Awareness in Cooperative Robotic Teams,* IEEE 0-8186-7108-4/95.

Pasteels, J. M., Deneubourg, J. and Goss, S. 1987. *Self-organization mechanisms in ant societies: Trail recruitement to newly discovered food sources,* in Pasteels, J. M. and Deneubourg, J. (eds.), "From individual to collective behavior in social insects", Basel: Birkäuser Verlag.

Pfeifer, Rolf 1995. *Cognition - Perspectives from autonomous agents,* Robotics and Autonomous Systems 15, pp47-70.

Pylyshym, Z. W. (ed.) 1987. *The Robot's Dilemma. The Frame Problem in Artifical Intelligence.*

Robins, R. H. 1997. *A Short History of Linguistics,* 4th edition (Longman Linguistics Library), Addison-Wesley Pub Co*;* ISBN: 0582249945.

Shepard, R. N. and Cooper, L. A. 1982. *Mental images and their transformations,* Cambridge: MIT Press/Bradford.

Stains, H. J. 1984. *Carnivores*, in Anderson, S. and Jones, J. K. Jr. (eds.), "Orders and Families of Recent Mammals of the World", John Wiley and Sons, N.Y., pp491-521.

Steels, Luc 1996. *The origins of intelligence,* Proceedings of the Carlo Erba Foundation, Meeting on Artificial Life. Fondazione Carlo Erba. Milano.

Suchman, L. A. 1987. *Plans and Situated Actions: The Problem of Human-Machine Communication,* Cambridge: Cambridge Press.

Thelen, Esther and Smith, Linda B. 1994. *A Dynamic Systems Approach to the Development of Cognition and Action,* A Bradford book, MIT Press, *ISBN 0-262-20095-3.*

Wilson, E. O. 1971. *The Insect Societies: Their Origin and Evolution,* New York: Narcourt, Brace & Co.

Wilson, E. O. 1975. *Sociobiology: The New Synthesis,* Harvard.

Yuta, S., Suzuki, S. and Iida, S. 1991. *Implementation of a small size experimental self-contained autonomous robot - sensors, vehicle control, and description of sensor based behavior,* Proc. Experimental Robotics, Tolouse, France, LAAS/CNRS.

Zelinsky, A., Kuniyoshi, Y. and Tsukue, H. 1993. *A Qualitative Approach to Achieving Robust Performance by a Mobile Agent,* Robotics Society of Japan Conference, Japan.

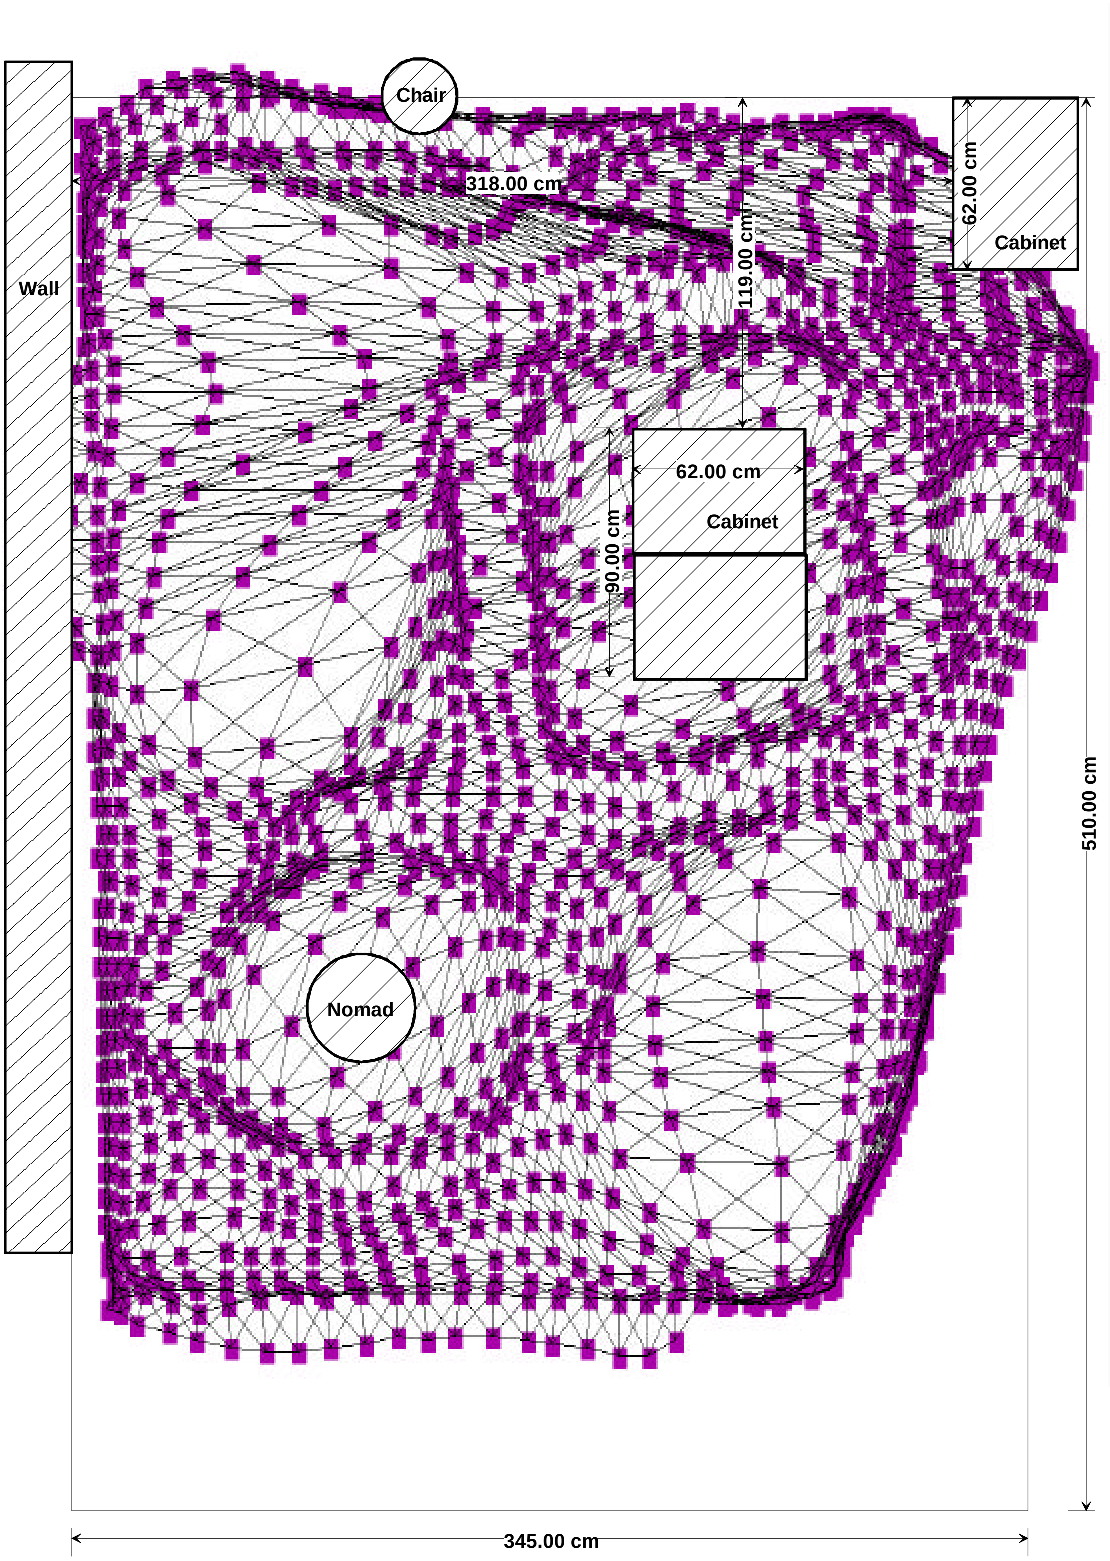

Figure 16 - SOM generated by Flo over part of our Lab