

# A Local Approach to Developing Grounded Spatial References in Multi-Robot Systems

Nathan Wiebe and John Anderson  
Autonomous Agents Laboratory  
Department of Computer Science  
University of Manitoba  
Winnipeg, Manitoba, Canada R3T2N2  
andersj@cs.umanitoba.ca

**Abstract**—For a mobile robot to be able to communicate usefully with others, the symbols it uses to communicate must be grounded to entities in the environment, and those groundings made consistent among agents. While it is common practice to hand-construct such groundings, this does not scale to large problems. In particular, when communicating about useful spatial references, there are a large number of potentially relevant groundings, even for a basic task such as navigation. This paper describes the development and evaluation of an approach that allows a group of robotic agents to develop consistent shared groundings for locations in an environment over time. This approach is based on local communication and interaction, and does not rely on the ability to broadcast references to all agents, and so is suitable for domains in which communication may be sporadic, such as robotic rescue. The evaluation of this approach, which compares several different grounding techniques, shows that shared groundings can be developed effectively over time, and that these improve the effectiveness of communication in a multi-robot setting.

## I. INTRODUCTION: GROUNDED COMMUNICATION IN MULTI-ROBOT SYSTEMS

In order to realize the benefits of deploying a team of robots, individuals are often required to communicate to coordinate their activities. While the benefits of communication to everyday human cooperative activities are easily observed, communication has also been empirically shown to be of great benefit when solving problems in multi-agent systems. Matarić [11] showed that communication was necessary to deal with credit assignment problems in multi-agent learning, for example, while Balch and Arkin [1] showed the advantages of employing communication in team-based foraging tasks. Even in domains such as robotic rescue, where interference and disruption of infrastructure may limit communication to short ranges or line-of-sight, communication can still be effective in improving the group navigation and mapping of a robot team.

Communication, however, can only be of use to the degree that the symbols used for communication are associated with (*grounded* to) concepts in the environment consistently among communicating agents. The problem of creating and maintaining these groundings is known as the *symbol grounding problem* [9]. In most multi-agent

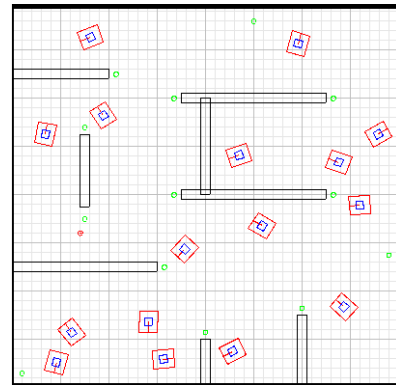


Fig. 1. 16 robots in an 8m x 8m domain

domains, groundings across a population of agents are hand-constructed. While this is certainly possible to do for small environments, it becomes unwieldy very quickly as the complexity of the environment increases, and simply will not scale to environments of any reasonable size [10]. Such hand-constructed groundings are also biased to the perspective of the humans creating them, as opposed to what might be arrived at by the agents in the domain themselves [10].

The problem of scale in providing groundings becomes much more significant when agents are mobile. In mobile settings, robotic agents must typically communicate spatial information, and thus have potentially a very large number of grounded physical locations that they might wish to communicate to others. Typically a shared coordinate system and the ability to localize well within it are provided by system designers, so that all references to locations become absolute. Robots entering a collapsed building from different starting points, for example, might share a common GPS coordinate system and know one another's starting location to exchange partial maps in a setting such as that shown in Fig. 1.

When such a shared coordinate system is unavailable, being able to share references to spatial locations in a sensible way becomes a very difficult problem. If GPS were unavailable, for example, which could happen due to interference (e.g. in a disaster setting), locale (e.g. underground) or nonex-

istence (e.g. on other planets), each robotic agent would have to begin with its own coordinate system, and any shared references would have to be built up across a population over time, so that an agent could refer to a location by name and have others understand it in terms of their own internal coordinate systems. The bounded resources at each agent's disposal preclude simply sharing each and every grounding among all agents - communication to agree on consistent symbol grounding in any reasonably complex domain would in many cases be combinatorially more significant than the amount of communication necessary for problem-solving itself. However, building up common-sense shared references to locations (and other entities) over time is a much more human-like approach to the problem, and one that serves as a more generally intelligent solution. A team of robots using such an approach would be able to function anywhere, and would be able to be developed independently of one another and form teams in an *ad hoc* manner.

This paper describes the development and evaluation of a multi-robot system which develops consistent shared groundings for locations in an environment across a population of agents without sharing a coordinate system, allowing agents to communicate information relative to those groundings. A location grounding, for our purposes, is a symbol that is associated with an agent's internal spatial coordinates, allowing the agent to make reference to a place in the environment (and ultimately be able to go there or refer to it in reasoning or communication). A shared grounding is one where the same symbol refers to the same absolute location in the world, even though the agents sharing the grounding may have completely different representations of that location internally. The approach presented here begins with no prior knowledge of the environment or shared points of reference, requires no explicit location labelling phase, requires only local encounters between agents rather than broadcast communication, does not flood the environment with groundings, and is demonstrated to be functional with larger populations (up to 16 robots) than previous work (e.g. [10], which uses only 2 robots). This approach is then evaluated in a domain where communication about locations in the environment can improve performance, and is shown to provide significant benefits, indicating that agents using this approach can reap the benefits of grounded communication without predefined groundings.

There are a number of ways in which references can be made to entities in the environment in a grounded fashion [7]. An *iconic reference* is an association based on observable features (e.g. seeing a fire and associating that sensor data with that concept). An *indexical reference* associates two icons (e.g. smoke may be an index for fire), while the most powerful is a *symbolic reference*, which is an arbitrary relationship between icons, indices, and/or other symbols. While associating sensory phenomena to internal representations in an iconic manner is a relatively straightforward form of learning, the other two types require much more sophistication. The approach described here demonstrates the use of both indexical and symbolic references.

The remainder of this paper overviews related work in this area, describes our approach and its implementation, and then describes the results of an empirical evaluation.

## II. RELATED WORK

Much practical reasoning employs symbols as placeholders for objects in the physical world. The groundings between perceivable objects and the symbols used to refer to them are commonly referred to as *anchors*, and the problem of developing and maintaining anchors for individual mobile agents through perception (a subset of the broader symbol grounding problem) is known as the *anchoring problem*. Coradeschi and Saffiotti presented a preliminary formal logic solution to the anchoring problem in 2000 [4], defining the major components in predicate logic. They later extended this work to address using symbols in actions and plans [5]. While this work provides a formal framework for maintaining anchors to symbols via perception in single agent systems, it does not deal with the multitude of practical issues that arrive when the shared groundings necessary for communication in multi-agent systems are considered, since the primary problems in group are not only deciding whether something is worth grounding, but reconciling the multitude of inconsistencies among the groundings of various agents.

Comparatively little work in grounding has been performed in multi-agent systems, because of the problems of developing and maintaining consistency of groundings among agents. Steels has done work in evolutionary linguistics, examining how individual agents in a group may be able to generate discrimination trees to distinguish one object from other in the environment, without the aid of a teacher [12]. Vogt [13] later improved on this work. These works are more limited than the work we present here, in that they operate with only a very limited number of objects in the environment, and the domain itself exists only to perform groundings as opposed to having groundings acquired over the course of activity. However, these do show that it is possible to allow agents to develop shared symbol associations in coherent ways.

Billard and Dautenhahn [3] studied the benefits of social skills to learning in heterogeneous multi-agent systems. In their environment, robotic agents were provided with a following behaviour, allowing them to keep a similar sensory context to another robot by maintaining physical proximity, and learned to associate words with color descriptions of flooring after receiving demonstrations by following a teacher broadcasting the correct color name. While agents could eventually become teachers, the language used never changed: groundings were never invented, merely spread from an individual teacher through the population. In the approach described in this paper, all agents are both teachers and learners all of the time, since each is developing useful groundings in the course of its own work, as well as trying to spread those groundings to others in order to be able to communicate about those spatial locations. This is thus suited to unknown environments where an existing competent teacher cannot be assumed.

Jung and Zelinsky [10] describe the implementation of a heterogeneous cooperative robotic cleaning task which benefits from the use of a symbolic references for communication. Two robots were required to clean a laboratory floor: *Flo* sweeps litter into piles, while *Joh* vacuums the piles to remove them. Their system adds three layers of behaviours for coordination between agents, each of which improves performance as it is added. First, Joh is given the ability to track Flo visually (useful here since Flo will likely be near a pile of litter). Next, Flo is given the ability to signal the location of a new litter pile to Joh, in coordinates relative to itself. Joh and Flo use the same wheel encoders, providing a pre-defined shared grounding for the communicated coordinates. Finally, communication using predefined symbolic references is employed. Each reference is communicated in terms of the angle and distance from a reference line connecting two points in the robots' environment, thus allowing a new point to be referred to using two points that are already known. This requires that robots go through an explicit location labelling phase in order to develop a common set of reference points so that communication can be understood. While this work shows that a human-centered coordinate system is unnecessary for navigation and communicating about locations, and that symbolic references are very useful for coordinated activity that takes place across a spatial area, the approach employed is not practical for problems of any significant size. Labelling locations is done by arbitrarily labelling enormous numbers of points (in excess of 500 for a 3.4 x 5.1 m room) during a phase in which both robots devote their full effort to collectively labelling locations (i.e. there is never a situation where one robot grounds a location independent of the other and then shares that grounding later). The number of points grounded is impractical for any everyday activity, but more importantly in a multi-robot setting, is impractical to make consistent across any significant population. This is unrealistic compared to most situations, where individuals are likely to take note of points that are interesting to them during the course of useful work, and share these with others when it is convenient, as is done in the scheme we present in the remainder of this paper.

### III. DEVELOPING GROUNDED SPATIAL REFERENCES

Any scheme for developing grounded spatial references among a team of agents requires three distinct policies: deciding when a particular location is worth grounding by an agent (i.e. allocating some label and maintaining an association for the agent's own activities); deciding when such a grounding is worth sharing with others; and deciding how to reconcile groundings when they conflict with those shared by others. We present these in the separate subsections that follow, with specific implementation details provided in Section IV. In our implementation, labels are random integers and individual coordinate systems are cartesian planes, but these techniques place no restrictions on the representation used for labels or coordinate systems, and no shared representation for coordinates is assumed.

#### A. Grounding Locations

Whether any given location in an environment is worth maintaining a correspondence with a label is largely dependent on the task a robot is performing. If one is cleaning a room, one will consider different locations to be of significance than if one is painting it. During these and many other tasks involving mobility, however, a robot is navigating the room, and so we attempt to define several approaches that will suit this task and thus indirectly assist in others. The utility of these schemes is examined in Section V.

The first of the strategies we employ, *label-at-meeting*, creates a shared grounding between a *pair* of agents at the current location when they encounter one another in the environment. This is a basic strategy that takes advantage of chance encounters that will inevitably occur in multi-robot domains. The dynamics of this strategy result in more groundings in areas where robots tend to spend most of their time, since encounters are more common in these areas. When one robot encounters another, it stops and invites the other to ground a symbol. If the other is amenable (since it may be busy with other things), it requests a name for the robot's current location. The original robot names the location and the new name is subsequently used by both. A delay is built into the implementation (Section IV) to stop this from occurring too frequently so agents are not taken away from their primary task in this environment.

While *label-at-meeting* is domain independent, it is also limiting in that an agent is restricted to creating a grounding only when another agent is present. The second strategy we consider, *label-spatial-entropy*, is more general in that it requires only a single agent and tries to predict, in a domain-independent way, which locations in the environment are likely to be useful. To make this prediction, we note the previous use of information theory in categorizing space for path planning. Baltes and Anderson [2] define the entropy of a spatial area (referred to as *spatial entropy* in this work), a metric that uses information theory to measure how mixed a given portion of the environment is in terms of free and occupied space. Spatial entropy is highest where there is the largest mixture of open and blocked areas (i.e. where the environment contains the most information). Here we employ this measure to choose useful locations for grounding: when an entropy threshold is exceeded, a grounding is generated (Section IV).

To contrast these general techniques and examine the effect of providing more specific information to agents ahead of time, we also define the *label-environment-feature* technique, where specific phenomena useful to the task at hand are explicitly described to the agent and grounded when these are recognized. This is a much more domain-specific technique, intended to serve as an upper bound for the quality of groundings that can be made in a given domain.

#### B. Sharing Groundings

The *label-at-meeting* strategy is the only one of those above that contains some means of sharing a grounding

among agents, and this is limited in that only the two agreeing to make the grounding possess it. To share groundings among agents generally, some means beyond mere communication must be provided. Since each agent may employ its own internal coordinate system, communicating these coordinates would be meaningless to others. Our approach employs the physical demonstration of groundings, to allow other agents to ground a location by perceiving it themselves.

Demonstrations are initiated when agents encounter one another. If an agent perceives another within a specific distance, and a known grounding is also within a given distance, the agent possessing that grounding will offer to demonstrate it to the encountered agent. If the second agent accepts the offer, it will signal the offering agent and begin to follow that agent as it moves to the demonstrated location. When the first agent arrives at the location of the grounding, it will then send a message to the second indicating the label it employs to designate the location. If this location is novel to the second agent, that agent will also ground the location to that label, and will signal the first agent that it has done so. This offer-and-acknowledge approach allows either agent to avoid participating in the demonstration if other tasks take priority, and in turn support balancing the importance of the overall task of the team with the benefits that another shared grounding might bring.

### *C. Consistency Among a Group of Agents*

An agent may learn many groundings by demonstration, but since each agent is also learning from others and making groundings itself, it will encounter many inconsistencies as well. The most important part of creating a set of groundings among a team of agents is mediating these conflicts appropriately as they are discovered, in a decentralized manner in which the computation is entirely local. Some mechanism is needed so that when groundings are shared, any conflict in labels or locations will be more likely to be resolved toward greater consistency among team members as a whole, rather than away from it.

In our approach, this is done by maintaining a reference count for each grounding within each agent. This reference count indicates, to the best of that agent's knowledge, the number of agents in total that employ this same grounding. This is a heuristic value, because it is maintained based on local encounters rather than global communication: many of the agents with which it has been shared will have since gone on to demonstrate the grounding to others, or may have a conflicting grounding replace this information. When performing a demonstration as described above, the demonstrating agent will send its reference count to the second agent in addition to communicating the symbol used, and will update this count by one if the second agent accepts its demonstrated grounding. In the case of a conflict during a demonstration, reference counts of the symbols involved allow a local determination of which is likely to be more widespread, and thus a means of encouraging less widely used groundings to be discouraged and those more widely used to be propagated.

In practice, there are three types of conflicts to be resolved. The most obvious is that the second agent (the viewer) may know the location grounding already, but under a different symbol, either because it has learned it from someone else or developed it itself. In this case, the agent with the lowest reference count gives up its grounding and adopts that with the higher reference count, and each updates the count accordingly. As this reference count becomes higher, it will eventually propagate to other agents and replace competing labels for the same groundings across the population. In the case of equal reference counts, the grounding of the second agent (the viewer) will be used in both agents: this leaves the two agents consistent in terms of this grounding after an encounter, and serves to add some change to the population and eventually cause one grounding to dominate the other.

The second type of conflict occurs when the second agent recognizes the symbol communicated, but already uses that symbol for a different location grounding. Here, consistency demands that one of these groundings must ultimately be abandoned, and the reference count is used to preserve the more prevalent of the two. A grounding is lost to one agent, but that grounding is less valuable to the group of agents as a whole than the one that is preserved. If the lost grounding is something of use, it will be rediscovered or relearned from others. It may be encountered again with the same symbol, but eventually one of the two competing uses will prevail across the population due to reference counts. The alternative to this, keeping the grounding and choosing yet another new name, is not followed because while it would preserve the grounding in one agent, it would cause even more groundings with redundant names to ultimately spread across the population, further reducing consistency. Each of these forms of conflict resolution requires a constant ( $\epsilon$ ) in the system to determine at what point two locations are considered far enough away from each other to be different: this is left to the particular implementation, since it depends on the size of the area involved and the size of the robots.

In the final form of conflict, both the location and the label are known to the viewing agent. This may in fact be the same grounding, but recorded slightly differently due to errors in perception or odometry. The  $\epsilon$  constant is used as a threshold to determine if two locations are the same, and if this is the case the higher reference count for this grounding will be used by both agents. If these locations are not the same, then the viewing agent has both a different symbol grounded to the current location, and the symbol used by the demonstrator in use for another grounding. This amounts to both the two prior cases, and these are resolved as described above.

## IV. IMPLEMENTATION

Since this approach is designed to operate as a group of robots performs some other duty, it could be implemented in any environment. To keep this general, we chose a basic multi-robot exploration domain, and implemented this approach in Java using Player/Stage [8]. Agents map the domain individually as they explore it, and can plot paths to

specific locations based on their map. The task environments were 11x11m and 8x8m arenas, containing walls, doors, and obstacles (e.g. Fig. 1).

Grounded communication in this domain is employed in the form of goal communication. A goal is placed randomly in the environment for the agents to find. As each agent explores the environment, it creates its own groundings using the methodologies described in Section III and shares these as other agents are encountered during the course of exploration. When an agent finds the goal, its location is broadcast to all agents in a grounded manner (this is the only broadcast communication used, and is used as a means of evaluating the efficacy of groundings in this domain). The subset of agents with the shared grounding(s) necessary for the goal location to have meaning are then able to plan a path and navigate to the goal. As greater consistency develops among a team of agents over time, all agents in the group should be able to arrive at the goal location more and more quickly.

Two different techniques are used to describe the goal location. The first is simply broadcasting the name for the grounding the discovering agent has nearest to the goal (an indexical reference). This location will not necessarily be at the goal itself, but it is the closest location to it that can be understood by others directly.

A second technique implemented was to specify the goal location using two grounded locations and a more complex symbolic reference. The symbolic reference is specified as  $(p_1, p_2, a, d)$ . The indicated position is found by drawing a reference line from  $p_1$  to  $p_2$  and then, starting at  $p_2$ , travelling  $d$  times the distance from  $p_1$  to  $p_2$ , at angle  $a \times 360^\circ$  to the reference line. This is similar to the symbolic reference used in [10]. Understanding these two grounded locations allows the basis of a shared coordinate system between communicator and receiver, and allows an agent to specify a goal location precisely even if it has no grounding near the actual goal. This flexibility comes at the cost of requiring two shared groundings instead of one.

Robot control is behaviour-based, and robots are given basic navigation and exploration behaviours (e.g. wandering, obstacle avoidance) and communication behaviors. Higher level behaviours support mapping and path planning. Mapping uses an occupancy grid approach, with each robot beginning at a 0,0 coordinate at its own starting location and making no assumptions about the size of the environment (thus requiring internal mapping to dynamically resize its world map as new areas are discovered). Path planning follows a quad-tree-based approach, with path improvements based on the technique of [6].

Behaviours are also implemented to encompass the three grounding generation strategies described in the previous section. For our implementation, label-spatial-entropy scans the fifteen occupancy grid cells (each cell is 10x10cm) centered around the agent every ten seconds to calculate spatial entropy, and the threshold for grounding was set to 0.75. To deal with domain-specific features for the label-environment-feature strategy while discounting errors in per-

ception that are not relevant to grounding per se, in our implementation we implemented feature labelling by placing perceivable feature markers that could be easily recognized. In all labelling strategies, we avoid flooding the environment with groundings by removing any groundings with a reference count of 1 within 1.5m whenever a new grounding is produced. Redundant groundings serve no purpose and since they have not been shared their removal has no impact on the rest of the population. All grounding strategies label new groundings with a randomly generated integer in the range 0 to 10000. If the new randomly generated name is already in use to ground a location within the agent, a new label is selected.

These grounding behaviours are then supplemented with the demonstration behaviour (with a 30-second delay forced between grounding demonstrations to allow agents to move away from an encounter without immediately initiating another) and the conflict resolution strategy. These encompass a number of lower-level behaviours (e.g. a behaviour to move to a nearby grounding to demonstrate it, follow another robot, etc.). Finally, a goal seeking behavior is implemented that attempts to perceive if the robot has reached a given goal, and listens to goal locations broadcast in terms of grounding labels by other robots. Further details on all of these behaviours are available in [14].

## V. EVALUATION

The evaluation of this approach was done in simulation using Stage [8] for a number of reasons. The primary reason was working in environments with larger populations of robots than were available physically. Issues of control and repeatability were also significant, since this evaluation involved repeatedly resetting a team of robots in an environment across many trials. In addition, Player/Stage has been validated as accurately simulating the behaviour of Pioneer robots, and code developed under Player/Stage will run under Pioneer robots directly.

The evaluation performed involved examining the efficacy of the grounding schemes described in Section III, in the domain described in Section IV. Agents each start in the domain with no groundings and a random placement and heading, and construct, demonstrate, and resolves conflicts between groundings as described previously. Whenever the goal is found by an agent, the location is broadcast using an indexical or symbolic reference as described in Section IV, and others understanding the groundings used now have the ability to find the goal more quickly. If an agent has not found the goal after 10 minutes of searching, it gives up. After all agents have found the goal or given up, one iteration of a trial is complete. During this iteration, agents collectively created a set of groundings which should be useful to improve their performance in this environment in future. These groundings should become more extensive and more consistent as agents gather more experience in this environment. In order to examine this effect, a full experiment consists of a series of 200 of these iterations, where agents maintain their groundings




Fig. 2. The split and hallway 8x8m environments.

between iterations and are placed in new random locations (with a new, randomly located goal) for each iteration.

Experimental trials were conducted in many different configurations. The three grounding creation strategies were explored independently, as were the two goal communication mechanisms. The size of team was varied (2, 4, 8, and 16 robots), as was the size of the domain (8x8m and 11x11m). We also attempted to vary the configuration of the space, since domains of a given size can have very different complexity, by defining four different layouts: open space for a baseline, a general partitioned domain meant to be typical of an office (Fig. 1), a *split* domain (Fig. 2) where barriers more strongly restrict travel from one end to the other, and a *hallway* domain with areas partitioned off either side. The distance at which a demonstration would be performed and the  $\epsilon$  constant were 5m and 50cm respectively (additional experiments also examined varying these [14]).

These experiments all tracked group performance over trials as shared groundings were developed, and also examined both the final density of groundings produced and the percentage of consistency in groundings among agents that was ultimately achieved. The latter was computed by gathering all individual groundings after an experiment into a full set, removing duplicate names for the same locations (where two robot widths - 80cm - was used as threshold for equality) and keeping only the most commonly used as the set of globally ground locations. The percentage in each agent that was consistent with this global grounding set was tallied, and these were averaged across all agents to produce an average consistency. This metric is a pessimistic one, in that by throwing multiple values away, we are discounting the partial consistency that does exist in the removed values (e.g. 49% of agents may have a consistent label for a location that is still useful to this subset, but this is removed in favour of the 51% that have a different label). This measure is also biased against larger populations, because unique individual groundings affect the measure more in larger populations. For example, if there is only a single globally ground point, and one agent that possesses that grounding, this amounts to a 50% global consistency in a 2-robot team, but only a 6.25% consistency in a 16-robot team.

In all cases, results are presented by comparing the team performance against a control group consisting of the same configuration but without the ability to develop and communicate with shared groundings, and expressed as a percentage of improvement over this baseline. Here we present a representative subset of the results obtained,

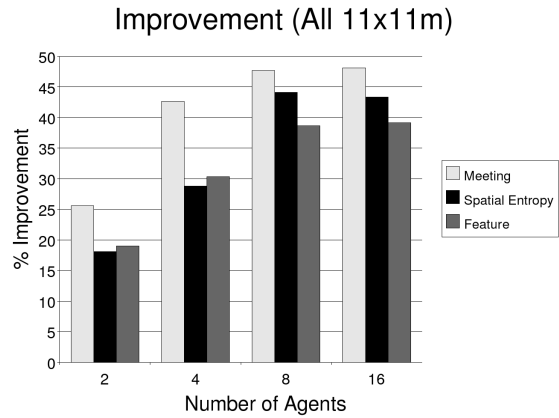


Fig. 3. Average percentage improvement by grounding strategy in 11x11m environments.

focusing only the larger environment size, and the reader is referred to [14] for further results of interest.

Over all domain types, in the larger domain size, the average improvement using a goal communicated as a single grounding (an indexical reference) is shown in Fig. 3. This showed good improvement as team sizes increased up to 8 robots, followed by similar performance when doubling the population to 16 robots. In an 8x8m domain, the performance of a 16-agent group dropped, and in both sizes this is showing the point at which overcrowding and the resulting interference between agents (e.g. collisions, more time spent avoiding obstacles) tended to limit the functionality of a larger team. Nonetheless, the techniques described here produced observable performance gain.

The label-at-meeting strategy had the largest performance improvement. This is attributed to already having a shared grounding between two agents when the grounding is made, leading to an immediate increase in the ability to spread that grounding quickly. This is especially significant for 2-agent teams, since a shared grounding is already globally consistent. Looking at consistency (Fig. 4) in these environments illustrates this. In this task, meetings were naturally distributed across the environment, leading label-at-meeting to produce wide coverage in groundings locations that led to a greater likelihood that a goal could be usefully described to most agents. In situations where meetings were infrequent (e.g. a small team over a large domain), this would be less useful, but in such a situation there is also a much stronger likelihood of agents working independently in different areas, and arguably less of a need for grounded spatial communication than in closer confines.

Label-spatial-entropy was generally the second most successful approach to choosing groundings. This was despite the fact that it produced significantly more groundings than the other two strategies, especially with large numbers of agents (Fig. 5). This increase is still not large, and the number of points grounded by 16 robots is far fewer than that of [10] makes for 2 robots. For 2 and 4 agent populations, agents are able to exchange enough locations to maintain a reasonable

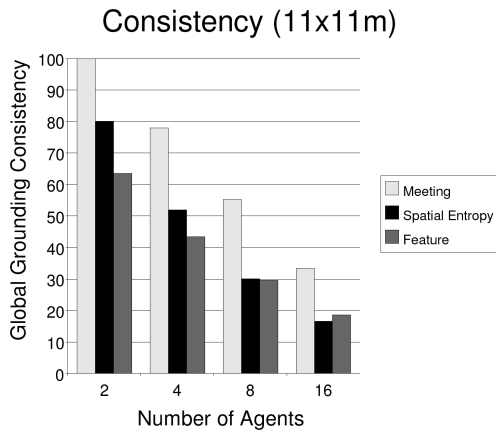


Fig. 4. Average global grounding consistency by grounding strategy (indexical reference) averaged over all 11x11m environments.




Fig. 5. Grounding density by grounding strategy (indexical reference) averaged over all 11x11m environments.

level of consistency. With larger population sizes, the number of groundings exceeds the number that can be reasonably shared, leading to lower consistency. This does not rule out spatial entropy as an valuable alternative: a domain could be adjusted to use a higher entropy threshold and still use this reasonably domain-independent technique.

It may seem surprising that grounding random encounters performed better than a technique that was specifically designed to ensure useful areas for navigation were grounded. However, not only is there greater sharing built into the label-at-meeting strategy, the points being grounded by label-spatial-entropy are not being exploited in this domain as strongly as they might be. Despite these groundings being interesting from the standpoint of navigation, the grounded communication employed references only the goal location. In that respect, these points are no better than other random locations. We anticipate that if an environment were to be set up where the grounded communication was more navigationally-oriented, spatial entropy would be more successful. An example of this would be an environment where information about waypoints to a goal were communicated instead of the goal itself.

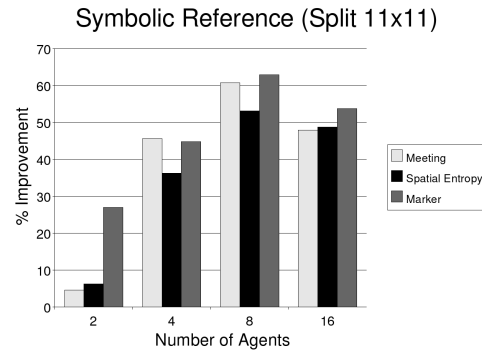


Fig. 6. Average improvement by grounding strategy (symbolic reference) in the split 11x11m environment.

Labelling environment features produced the lowest overall improvement, mainly because there were too few locations to be grounded (as can be seen from the grounding density). Having few locations to ground seems like a factor that should encourage high consistency. However, it also limits the opportunities for an agent to demonstrate an already-grounded location to others, since these locations are only demonstrated when agents encounter one another with a grounding nearby.

Because of the large numbers of trials demanded by the breakdown of all the controlled variables being evaluated, and the time taken to run trials (One trial of 200 iterations with 16 agents in one environment, for example, took six days using 2 computers), we examined some factors only in subsets of these environments. In particular, the effect of using symbolic references involving two groundings rather than using a single grounding as an indexical reference, was examined only in the split environment configuration. This was selected because it showed a high average improvement in the baseline experiments, and it also had high absolute times, and so we should expect a more sophisticated technique to be just as well-received. Results with goals communicated to team members symbolically using two grounded locations as described in Section IV in an 11x11m split configuration are illustrated in Fig. 6. Because this affects only how goals are communicated, density and consistency information should not differ the results presented earlier. To provide a baseline for comparison, the results shown in Fig. 3, disaggregated to show only the split domain, are shown in Fig. 7.

From these Figures, it can be seen that the label-at-meeting grounding strategy using an indexical reference to communicate a goal performs as well or better than when using a symbolic reference. The label-at-meeting environment provides good distribution of shared groundings over the environment, with high consistency, allowing the goal's location to be specified accurately enough with the single reference. This is coupled with the fact that a symbolic reference requires twice as many shared groundings for the reference to be useful, which leads to approximately equal performance between the two references.

The label-spatial-entropy grounding strategy shows little




Fig. 7. Average improvement by grounding strategy (indexical reference) in the split 11x11m environment only.

difference in performance across all agent populations, using either form of goal reference. The increased flexibility of a symbolic reference seems to be mitigated by the need to share twice the number of groundings in order to make use of it. Future work could explore this by selecting a domain in which a simple reference is less likely to be useful.

The label-environment-feature grounding strategy fared better than all others for an 11x11m split environment. We believe this is due to the advantage of symbolic references. Since agents using the label-environment-feature grounding strategy are limited to grounding certain locations, a single grounded reference has limited ability to specify the goal's location. This is in contrast to a more complex symbolic reference, which gives label-environment-feature the ability to specify an arbitrary point in the environment, and hence the goal's location, more accurately.

## VI. CONCLUSIONS

In this paper, we have presented an approach to developing consistent shared groundings over time that uses local interaction between agents to bring consistency to a population. We have also described an implementation and evaluation of that approach. In comparing techniques for grounding, label-at-meeting was the most successful, because the implicit sharing when agents create a new grounded location leads to a high global location consistency, which improves the effectiveness of communication about the goal. This approach serves to allow agents to use their own independent spatial representations and coordinate systems internally, and develop the means to function together as a team without assuming a shared coordinate system or external facilities such as GPS. This in turn leads to more general agents. It also allows groundings to be developed that are naturally useful from the point of view of the agents themselves, rather than the viewpoint of the agent developers. Such non-anthropocentric mechanisms are an important part of developing truly autonomous agents that can function generally across a wide range of environments.

In this approach, the symbols used are initially undefined, as opposed to known by an initial teacher, and assumes no fixed reference points or common coordinate system, as opposed to the approach presented in [3]. Unlike the

approach presented in [10], this approach is carried out during other work as opposed to during a labelling phase, grounds a practical number of points, and is shown to work with greater than 2 agents.

We are currently exploring a number of avenues of future research. While we have explored the efficacy of different grounding strategies, this work has not yet tried to compare different strategies for resolving conflicts between agents. At a higher level, some important areas that we are exploring include actively balancing grounding demonstrations with the work already being performed in the domain, and attempting to recognize when sufficient groundings have been made. While the latter is very difficult to recognize from a global perspective, it is conceivable that useful heuristics may approximate this and be general enough to be useful in many domains. We are also exploring issues of perception of locations worth grounding, through vision and other means, with real pioneer robots, to better integrate this with the other fundamental problems that robots must solve.

As a final note, it is noted in [14] that a bug in recognizing the second category of conflict detailed in Section III-C may have negatively influenced the results obtained. These experiments were entirely re-done with this bug fixed, and resulted in only a very small (1%) difference.

## REFERENCES

- [1] T. Balch and R. Arkin. Communication in reactive multiagent robotic systems. *Autonomous Robots*, 1(1):27–52, 1994.
- [2] J. Baltes and J. Anderson. Flexible binary space partitioning for robotic rescue. In *Proceedings of IROS-03*, pages 3144–3149, Las Vegas, 2003.
- [3] A. Billard and K. Dautenhahn. Experiments in learning by imitation – grounding and use of communication in robotic agents. *Adaptive Behavior*, 7(3/4):415–438, 1999.
- [4] S. Coradeschi and A. Saffiotti. Anchoring symbols to sensor data: preliminary report. In *Proceedings of AAAI-00*, pages 129–135, Austin, 2000.
- [5] S. Coradeschi and A. Saffiotti. An introduction to the anchoring problem. *Robotics and Autonomous Systems*, 43(2-3):85–96, 2003.
- [6] I. Davis. Warp speed: Path planning for star trek: Armada. In *Proceedings of the AAAI Spring Symposium on AI and Interactive Entertainment*, 2000.
- [7] T. Deacon. *The Symbolic Species: the co-evolution of language and the human brain*. Penguin Books, Toronto, 1997.
- [8] B. Gerkey, R. Vaughan, K. Støy, A. Howard, G. Sukhtame, and M. Mataric. Most Valuable Player: A Robot Device Server for Distributed Control. In *Proceedings of IROS-01*, pages 1226–1231, Wailea, Hawaii, 2001.
- [9] S. Harnad. The symbol grounding problem. *Physica D*, 42:335–346, 1990.
- [10] D. Jung and A. Zelinsky. Grounded symbolic communication between heterogeneous cooperating robots. *Autonomous Robots*, 8(3):269–292, 2000.
- [11] M. Mataric. Using communication to reduce locality in distributed multiagent learning. *Journal of Theoretical and Experimental AI*, 10(3):357–369, 1998.
- [12] L. Steels. Constructing and sharing perceptual distinctions. In M. van Someren and G. Widmer, editors, *Proceedings of the European Conference on Machine Learning*, pages 4–13. Springer-Verlag, 1997.
- [13] P. Vogt. Symbol grounding in communicative mobile robots. In *Proceedings of the AAAI Fall Symposium on Anchoring Symbols to Sensor Data*, 2001.
- [14] N. Wiebe. Developing grounded communication in multi-agent systems. Master's thesis, Department of Computer Science, University of Manitoba, Winnipeg, Canada, January 2006.