Accepted for publication in ICRA 2013.
Citation: Heath S., Ball D., Schulz R., and Wiles J., Communication between Lingodroids with different cognitive capabilities, Proceedings of the 2013 IEEE International Conference on Robotics and Automation, Karlsruhe, May 6 - 10, 2013.

Communication between Lingodroids with different cognitive capabilities

Scott Heath, David Ball, *Member, IEEE*, Ruth Schulz, *Member, IEEE*, Janet Wiles, *Member, IEEE*

Abstract—**Previous studies have shown how Lingodroids, language learning mobile robots, learn terms for space and time, connecting their personal maps of the world to a publically shared language. One caveat of previous studies was that the robots shared the same cognitive architecture, identical in all respects from sensors to mapping systems. In this paper we investigate the question of how terms for space can be developed between robots that have fundamentally different sensors and spatial representations. In the real world, communication needs to occur between agents that have different embodiment and cognitive capabilities, including different sensors, different representations of the world, and different species (including humans). The novel aspects of these studies is that one robot uses a forward facing camera to estimate appearance and uses a biologically inspired continuous attractor network to generate a topological map; the other robot uses a laser scanner to estimate range and uses a probabilistic filter approach to generate an occupancy grid. The robots hold conversations in different locations to establish a shared language. Despite their different ways of sensing and mapping the world, the robots are able to create coherent lexicons for the space around them.**

I. Introduction

A key challenge for performing useful tasks with a team of heterogeneous robots and humans is the ability to communicate effectively between agents with different embodiment and cognitive capabilities. The embodiment of an agent consists of its sensors, actuators and physical body, while cognitive capabilities include learning and language abilities. Lingodroids – language learning robots – have been used to model cognitive processes ranging from knowledge representation and planning to language development, symbol grounding and even imagination. A key aspect of lexicon learning is the connection of a word to its meaning, called "symbol grounding". The Lingodroids' advanced navigation skills have been particularly useful in learning lexicons that name spatial aspects of their environments.

Practical symbol grounding studies to date have rarely examined populations of heterogeneous robots. Robots with different sensors but identical cognitive representations were studied in [1]. It is still an open question what level of communication can be achieved between agents with different cognitive capabilities.

Previous work in the Lingodroids project has shown that real robots can learn a language for human-like concepts of space (locations and spatial relationships [2]), and that this framework can be extended to temporal concepts of durations [3]. In previous Lingodroid studies, the robots constructed their maps independently, and so have had unique cognitive maps of the world, but within each study the robots were functionally identical, using the same type of physical robot and the same underlying mapping algorithms and behaviors.

In this paper, we investigate Lingodroids with different embodiment and cognitive capabilities. Studies were performed with pairs of real robots. The robots were based on a rat-sized robot called the intelligent rat animat technology, or iRat, developed at the University of Queensland. While the iRats had the same physical size, actuators, and language systems, they differed in three key aspects:

1. **Sensors -** One robot used the iRat's single standard forward facing camera, which provided color images that are converted into appearances ('camera iRat'). The other robot instead used a 240 degree laser scanner that provided metric range information ('laser iRat').
2. **Algorithmic approach -** The camera iRat used the biologically-inspired RatSLAM system [4] which uses a continuous attractor network to filter appearances and self motion. The laser iRat used a probabilistic filter approach to localize where particles represent possible poses in the map [5].
3. **Spatial representations -** The camera iRat constructed and relaxed a semi-metric topological map, called the experience map. The laser iRat created occupancy grid maps offline from laser scan information, using particles to represent possible maps [6].

When agents have identical embodiment and cognitive capabilities, it is theoretically possible to transfer knowledge directly from one agent to the next. When such capacities differ, a direct transfer is no longer an option. Symbol grounding must be achieved through private grounding using different algorithms for each robot type, and through social grounding that is appropriate for all of the robot types.

In this paper we show that Lingodroids with different sensor types and map representations can develop coherent symbols for places, distances, and directions. Lingodroids is a good candidate for facilitating communication between teams of heterogeneous robots and humans, due to the human-like concepts learned in previous studies and its coupling with state-of-the-art SLAM systems [2]. The major contribution of this paper is the demonstration of spatial

S.H., R.S. and J.W. are with the School of ITEE at The University of Queensland, Brisbane, QLD, 4072, Australia, e-mail: scott.heath@uqconnect.uq.edu.au; ruth@itee.uq.edu.au; j.wiles@uq.edu.au.
D.B. is with the School of Electrical Engineering and Computer Science at the Queensland University of Technology, QLD, 4001, Australia. e-mail: david.ball@qut.edu.au

language learning on real robots with different embodiment and cognitive capabilities.

The paper presents a brief review of related work before providing details about the robot platform and algorithms for building maps and grounding language. We describe the experimental setup and present results that show the coherence of the spatial lexicons. The discussion focuses on the potential extensions of the methodology.

II. Literature Review

Heterogeneous robot teams are a growing research area, involving robots with a variety of abilities interacting, cooperating, coordinating, and communicating. Examples of tasks for teams of robots include environment mapping [7], cooperative localization [8], search and rescue [9], and decentralized environment modeling [10].

A key challenge for robots that are part of heterogeneous teams of robots and humans is how to communicate about information in their respective knowledge bases, formed through their individual interactions with the world. The shared language used for communication must be grounded in each robot's own representations, thus addressing the challenge of the symbol grounding problem referred to in the introduction [11]. To effectively communicate with each other, the robots need to link individual experiences with symbols via private or physical grounding [12], and develop a standard usage of shared terms via social grounding [13]. One solution to the symbol grounding problem involves robots learning categories embedded in the robot's sensorimotor interactions by playing language games [14]. Many variations on language games have been developed for different tasks and environments [13, 15, 16] including spatial locations and relations [1, 17-19].

Cognitive capabilities for mobile robots may differ in the mapping systems used. Simultaneous Localization and Mapping (SLAM) systems can vary in the sensors and the mapping algorithms used. Two distinct approaches to solving the problem of SLAM are probabilistic approaches [6, 20] and biologically inspired approaches [4, 21].

III. Method

This section describes and the robot's mapping systems. Note that the robots use the same Lingodroid communication system; however, the connections differ from the lexicons to the maps.

A. *Robot platforms and environment*

The same base robot platform, the iRat [22], is used for both agents. The iRat is the same size and mass as a large rodent, with an onboard 1GHz computer and wireless 802.11g/n. The robot moves about its environment using a differential drive system. Sharp infrared sensors orientated at -45, 0 and 45 degrees provide range information. The robot runs the Robot Operating System (ROS) [23] on Ubuntu and communicates to clients over wireless at 20Hz to send and receive desired and actual velocity commands.

For the studies in this paper the iRat uses its infrared range sensors for avoiding obstacles such as the walls and other iRats. The iRat attempts to wall follow down the center of a corridor. When it arrives at an intersection it randomly chooses a direction for exploration.

The two iRats have different primary sensors (see Fig. 1). The camera iRat is unmodified from the standard build and uses a forward facing widescreen camera with a horizontal field of view of 110 degrees. A ROS node publishes compressed 416x240 pixel color JPEG images from the robot's forward facing camera.

The second iRat, the laser iRat, has had the camera replaced with a Hokuyo URG laser sensor that is mounted upside-down so that it can sense the range to the walls of the environment. The laser scanner has a scan angle of 240 degrees with an angular resolution of 0.36 degrees and a detection range of between 20mm and 1000mm. A standard ROS node driver is used to publish the laser scans.

These studies were performed within an environment made specifically for the iRats, modeled on a map of Australia. An overhead view of the set, which measures 3.2 x 2.4 meters is shown in Fig. 2. Prominent Australian features such as Uluru (the red rock in the center) and the Sydney opera house (bottom right) can be seen.

B. *SLAM systems*

The robots' SLAM systems are completely different, one based on a biologically inspired topological approach, the other based on a probabilistic metric approach. Previous Lingodroid studies have only used the biological system.

The biological system is called RatSLAM and is inspired by the rodent hippocampus [4]. It has three parts: local views, pose network, and experience map relaxation. The local view part uses the camera images to match locations based on appearance similarity. The pose network is an energy based continuous attractor network which filters local view appearances and self motion estimates. The experience map builds and relaxes a semi-metric topological map. Although an initial map is created at the beginning of the study, the RatSLAM map is continuously corrected during usage.

The probabilistic system uses Gmapping [6] for map construction and Adaptive Monte-Carlo Localization (AMCL) [5] for robot localization within the map. Gmapping uses a particle filter to build occupancy grids from metric range and self motion information. Each particle carries an individual map of the environment, and the filter attempts to reduce the number of particles. The occupancy grid is constructed at the beginning of the study. The robot is then localized during studies using AMCL, a system which maintains a probability distribution using particles over possible robot poses within a static occupancy grid. The particle distribution is spread and moved based on robot motion and resampled based on laser range data. AMCL requires a relatively accurate estimate of the robot's motion, which is not met by the iRat's coarse wheel odometry.

Fig. 1. Lingodroid iRats in conversation. On the right is a standard 'camera iRat' which has a forward facing camera inside the dot of the 'i'. On the left is the modified 'laser iRat' which has a laser scanner mounted upside down. The standard iRat's cover has been removed to fit the laser scanner. The two iRats are shown in the environment used for the studies in the paper (also see Fig. 2).

Instead, local laser scan matching [24] is used to provide the motion estimate. Note that RatSLAM tolerates the inexact wheel odometry due to the use of a topological map.

To counter the uncertainty in localization using AMCL with the iRats, the laser iRat did not consider localization accurate if AMCL's global particle covariance was too high (determined by summing the covariance matrix elements and setting a threshold of 0.1 for these studies).

The important differences between the two mapping systems from a grounding perspective are the output representations. The output representation of RatSLAM is a topological map, expressed as a graph, whereas the output representation of Gmapping is an occupancy grid, expressed as a 2D array (see Fig 3). The output of Gmapping remains static after the initial map-creation phase, however, the output of RatSLAM continues to change and correct, as discussed further in the next section. Gmapping provides additional information about the location of obstacles compared with the map provided by RatSLAM. The occupancy grid produced by Gmapping will be 'metric' (distances and directions will be very accurate), while RatSLAM provides only a semi-metric map (distances and directions are only partially accurate).

C. *Language platform: Lingodroids*

Lingodroids develop lexicons using pairs of robots to evolve a shared language over a series of conversations, which are short social interactions [17]. The most basic conversations consist of a single question and response. The robots converse to learn words for locations called toponyms, which literally means 'place names'.

The robots use *where-are-we* conversations to create names for shared locations. Over many such conversations, the robots agree on a set of toponyms referring to different locations in the environment. This toponym lexicon can then be used to bootstrap generative conversations, such as

Fig. 2. The environment used for the studies in the paper, modeled after a map of Australia. The two iRats are shown interacting in the center right of the image.

how- far and *what-direction* – conversations that allow the Lingodroids to form a set of lengths and directions corresponding to the distances and angles between toponyms. This bootstrapping involves indirectly grounding the elements, a process known as 'grounding transfer' [25].

Words are associated with component parts (called 'concept elements') of the toponyms, distances, and directions. For the toponyms, different location-related concept elements were used for each SLAM system: for the camera iRat, toponyms were associated with RatSLAM experiences; for the laser iRat, toponyms were associated with grid square locations in an occupancy grid. RatSLAM experiences record positions in meters in abstract space while the grid squares are recorded in pixels. A key difference between the two types of concept elements (experiences vs pixels) is that RatSLAM experiences can move, as the map is continually corrected during usage. Having the concept elements move allows words to change meanings to correct associations formed when the robot was incorrectly localized. The occupancy grid created using Gmapping does not change during use, therefore words associated with that map never change meaning. For both systems, distance and direction concept elements are created as needed when referred to during a conversation. Distances are calculated from the estimated metric distance between locations in each robot's map. Directions are calculated from the angle between three locations in each robot's map.

Associations between words and concept elements are stored in *distributed lexicon tables.* Distributed lexicon tables maintain a set of all of the words used by either agent during conversations, a set of all the concept elements and a set of edges between words and concept elements that have been used together, called associations. This data structure allows storage of many-to-many relationships between

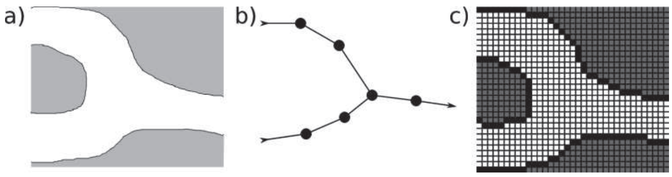

Fig. 3 RatSLAM and Gmapping use fundamentally different representations of their environment a) a junction where three paths meet b) the RatSLAM representation uses experiences connected by links. c) the Gmapping representation uses an occupancy grid where white indicates free space, black indicates an obstacle and grey indicates unknown.

words and concept elements. The strength of an association is the number of times that its <word, concept element> pair has been used together in conversation. The resulting structure allows a single trial to define a word, with additional trials refining the use of a word.

When choosing a word to describe a concept element, the speaker finds the word-concept element pair with the highest confidence. The confidence value, h_{ij}, for a word, j, and a concept element, i, is calculated by

$$h_{ij} = \frac{\sum_{k=1}^{X} a_{kj}(D - d_{ki})/D}{\sum_{m=1}^{N} a_{mj}} \tag{1}$$

where X is the number of concept elements within a neighborhood of size D of the current concept element, i; a_{ij} is the number of times that the concept element, i, and the word, j, have been used together; d_{ki} is the distance between concept element k and i; and N is the total number of concept elements. The distance between two concept elements is defined as the Euclidean distance between the coordinates of the elements.

The neighborhood size determines the coverage of the toponyms formed in each map: For toponyms and distances it was set to 0.3m for the camera iRat and 100px for the laser iRat, and for directions for both systems it was set to 60°.

Words are invented with probability, p, using the confidence value and a word invention temperature, as follows

$$p = k \exp\left(\frac{-h_{ij}}{(1-h_{ij})T}\right) \tag{2}$$

where $k = 1$, h_{ij} is the confidence value of the concept element-word combination, and T is the temperature, the word invention rate. A higher temperature causes words to be invented with higher frequency even when a valid generalization is available. T was set to 0.1 in this study for all concept types.

Lingodroids conversations are started by the robots whenever they have shared attention. Shared attention is established by an overhead camera, which detects when the two robots are within 50 pixels of each other, or 0.25m apart in the environment and notifies the two agents simultaneously with a boolean value.

Conversations using grounding transfer, such as *how-far* and *what-direction* do not depend on the robots' physical locations and so may be performed offline.

D. *Quality Measures*

Previous Lingodroid studies use coherence between lexicons as a quality measure. Coherence is calculated by rendering a lexicon onto a fixed resolution grid and then determining the number of matching grid squares as a percentage. However, in these studies it is not possible to calculate coherence of the toponymic lexicons directly, as the two mapping systems are not commensurate, in that features stored by one SLAM system have no representation in the other system.

The coherence of the toponymic lexicon was instead established by calculating the coherence of distance and direction lexicons that were constructed from it. Coherence was calculated for distance lexicons by choosing words to describe distances at 100 points from 0m up to the maximum distance in each robot's lexicon. The percent of matching words was then calculated between the two robots. For the direction lexicons, the word chosen was determined for every 2.5° from 0° to 360°. If the toponym lexicon is coherent and well grounded by each robot, then the *how-far* and *what-direction* conversations are likely to be coherent too.

IV. Experimental Setup

1. **Map building -** Both iRats independently explored the set and built their maps prior to starting the conversations. The camera iRat created a topological map using RatSLAM and the laser iRat created an occupancy grid map using Gmapping. The maps were saved at the end of this phase.
2. **Learning toponyms -** The two iRats were then placed back into the set together, the camera iRat running RatSLAM on the previously created topological map and the laser iRat running AMCL on the previously created occupancy grid. The iRats explored the set for two hours, localizing and holding *where-are-we* conversations when they moved within shared attention range. The conversations were used to create independent toponym lexicons.
3. **Learning distances and directions –** Distance and direction conversations were performed offline using the created maps and toponym lexicons from the previous phase to allow the creation of separate lexicons for distances and directions. 10 trials were done for this phase starting with the same initial maps and toponym lexicons from phase 2. 100 *how-far* and 100 *what-direction* conversations were held for each trial. Average coherences were calculated across the 10 trials.

V. Results

After their initial explorations in phase 1, each robot had individually mapped the area. After the *where-are-we* conversations, they had together constructed a shared toponymic lexicon (see Fig. 3). A total of 10 toponyms were created, with most (8/10) toponyms covering a contiguous region of the environment, and two toponyms (lenu and kumu) covering two local regions separated by an intervening toponym.

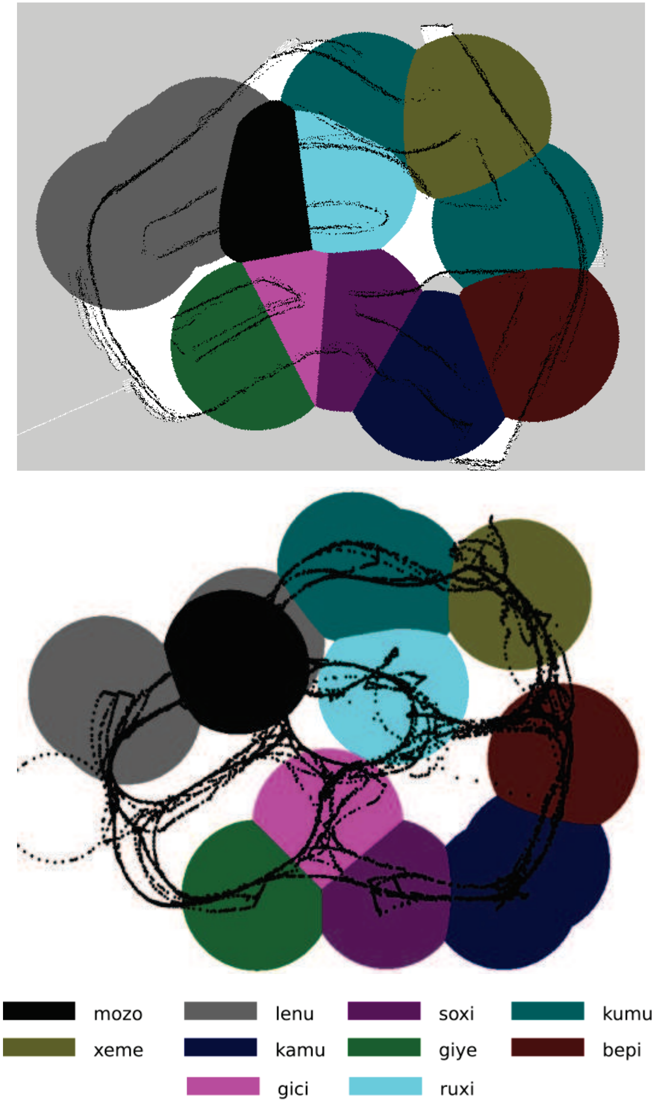

Fig. 4. Maps and Toponymic Lexicons developed for the laser iRat (top) and the camera iRat (bottom). The maps of both robots are recognizably the environment shown in Fig. 2, with a high degree of similarity between the locations for all of the toponyms.

Visual inspection indicates that the locations of toponyms in each map are similar; however, the maps from the two robots cannot be directly equated, since one is an occupancy grid and the other is a topological graph. What can be analyzed is the robots' functional use of the lexicon, by equating robot journeys on the maps. For example, to describe a journey starting in the north and following a clockwise journey around the outer perimeter of the set, each robot will pass through an almost identical sequence of terms. The edit distance between these two journeys is 2 (omission of kumu and addition of a second lenu in the camera iRat).

Following the 100 *how-far* and *what-direction* conversations, the robots had developed coherent distance and direction lexicons, with an average of 4.2 distance words and 4.5 direction words (averaged over 10 runs, see Table I, Fig. 4, and Fig. 5). The average coherence of the distance lexicons was 0.78 and the direction lexicons was 0.72.

TABLE I
DISTANCE AND DIRECTION

Measure	Distance (10 trials) $\mu(\sigma)$	Direction (10 trials) $\mu(\sigma)$
Coherence	0.78 (0.14)	0.72 (0.14)
Maximum Coherence	0.95	0.94
Number of Words	4.2 (1.1)	4.5 (0.9)
Number of Concept Elements	9.7 (0.88)	31.9 (2.4)

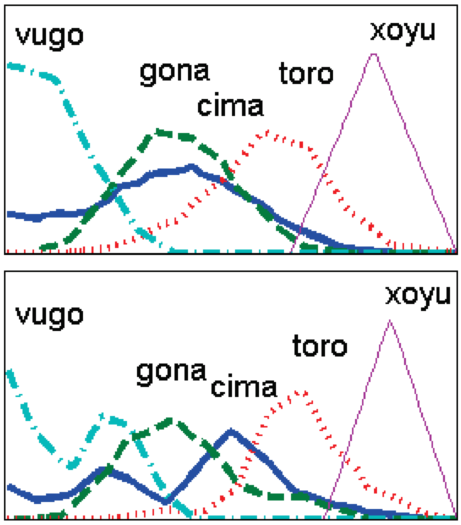

Fig. 5 The most coherent distance lexicon (coherence of 0.95) for the laser iRat (top) and the camera iRat (bottom).

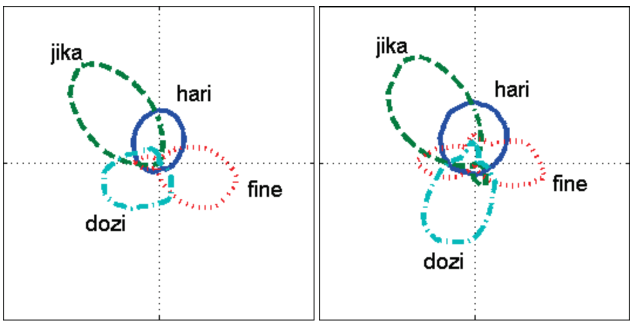

Fig. 6. The most coherent direction lexicon (coherence of 0.94) developed for the laser iRat (top) and the camera iRat (bottom).

VI. Discussion and Conclusions

These results show how communication can be achieved between agents with different sensors and mapping systems through shared externally grounded symbols. Experiences for each agent give rise to subjective characteristics that cannot be shared between them, but that does not preclude the evolution of a language for describing the world around them.

Toponyms learned by the Lingodroids are grounded in the very different representations of space formed from their characteristic sensors (laser vs vision). Features of these systems cannot be shared by direct transfer, nor can they interpret each other's maps. Instead, the two Lingodroids are able to ground their respective representations in shared experience. Each robot uses the shared experience to determine its own appropriate features and map.

It is likely that the more similar two agents are, the closer their subjective experiences will be [26]. However, even in previous Lingodroid studies with almost identical architectures [17], direct transfer would still fail due to the subtle differences between the robots' sensors.

There are limits to the differences between agents using the Lingodroids methodology. In common with other robot language studies, agents must share a common process for learning a shared symbol, including hearer and speaker roles and sharing attention [1] and each agent's referents must be reliable [1, 14, 16].

Granularity is an important consideration when dealing with continuous concepts, such as space or time [27]. An underlying assumption of the Lingodroids' studies is that the reference of a location term extends to a radius around specific points where the term has been used [2]. This simple version of location proves sufficient for bootstrapping distance and direction relations, and it is expected that with the same representations, the Lingodroids could be extended with minor modifications to learn spatial propositions typically studied in grid worlds [18, 19]. However, more complex concepts of location can refer to the space occupied by an object or some superset or subset of that space [27]. A constraint on using the Lingodroids methodology is that when attending a location the relevant spatial referent must be shared by both agents.

One of the fundamental assumptions in human language used to be that shared understanding must be grounded in shared biology. The Lingodroid studies reported here complement human-robot studies in demonstrating for the most fundamental of all lexicons – describing practical terms for space – it is possible bridge communication barriers across agents with different cognitive capabilities.

ACKNOWLEDGEMENTS

This research was supported by grants from the Australian Research Council and NSF Temporal Dynamics of Learning Center to JW.

REFERENCES

[1] D. Jung and A. Zelinsky, "Grounded symbolic communication between heterogeneous cooperating robots," *Autonomous Robots,* vol. 8, pp. 269-292, 2000.

[2] R. Schulz, G. Wyeth, and J. Wiles, "Lingodroids: socially grounding place names in privately grounded cognitive maps," *Adaptive Behavior,* vol. 19, pp. 409-424, 2011.

[3] S. Heath, R. Schulz, D. Ball, and J. Wiles, "Lingodroids: learning terms for time," presented at the The International Conference on Robotics and Automation, Saint Paul, Minnesota, USA, 2012.

[4] M. Milford and G. Wyeth, "Persistent Navigation and Mapping using a Biologically Inspired SLAM System," *The International Journal of Robotics Research,* vol. 29, pp. 1131-1153, August 1 2010.

[5] D. Fox, W. Burgard, F. Dellaert, and S. Thrun, "Monte Carlo Localization: Efficient Position Estimation for Mobile Robots," in *Proceedings of the Sixteenth National Conference on Artificial Intelligence (AAAI'99)*, ed, 1999.

[6] G. Grisetti, C. Stachniss, and W. Burgard, "Improved Techniques for Grid Mapping With Rao-Blackwellized Particle Filters," *IEEE Transactions on Robotics,* vol. 23, pp. 34-46, 2007.

[7] R. Simmons, D. Apfelbaum, D. Fox, R. P. Goldman, K. Z. Haigh, D. J. Musliner, M. Pelican, and S. Thrun, "Coordinated deployment of multiple, heterogeneous robots," 2000, pp. 2254-2260 vol. 3.

[8] L. E. Parker, B. Kannan, F. Tang, and M. Bailey, "Tightly-coupled navigation assistance in heterogeneous multi-robot teams," 2004, pp. 1016-1022 vol. 1.

[9] R. R. Murphy, J. Casper, M. Micire, and J. Hyams, "Mixed-initiative control of multiple heterogeneous robots for urban search and rescue," 2000.

[10] E. Gil Jones, B. Browning, M. B. Dias, B. Argall, M. Veloso, and A. Stentz, "Dynamically formed heterogeneous robot teams performing tightly-coordinated tasks," 2006, pp. 570-575.

[11] S. Harnad, "The symbol grounding problem," *Physica D: Nonlinear Phenomena,* vol. 42, pp. 335-346, 1990.

[12] R. A. Brooks, "Elephants don't play chess," *Robotics and Autonomous Systems,* vol. 6, pp. 3-15, 1990.

[13] A. Cangelosi, "The grounding and sharing of symbols," *Pragmatics and Cognition,* vol. 14, pp. 275-285, 2006.

[14] L. Steels, "Language games for autonomous robots," *IEEE Intelligent systems,* vol. 16, pp. 16-22, 2001.

[15] S. Kirby and J. Hurford, "The emergence of linguistic structure: An overview of the iterated learning model," *Simulating the evolution of language,* pp. 121-148, 2002.

[16] P. Vogt, "The physical symbol grounding problem," *Cognitive Systems Research,* vol. 3, pp. 429-457, 2002.

[17] R. Schulz, A. Glover, M. Milford, G. Wyeth, and J. Wiles, "Lingodroids: Studies in Spatial Cognition and Language," presented at the The International Conference on Robotics and Automation, Shanghai, China, 2011.

[18] L. Steels, "A self-organizing spatial vocabulary," *Artificial life,* vol. 2, pp. 319-332, 1995.

[19] A. Cangelosi, K. Coventry, R. Rajapakse, D. Joyce, A. Bacon, L. Richards, and S. Newstead, "Grounding language in perception: A connectionist model of spatial terms and vague quantifiers," *Progress in Neural Processing,* vol. 16, p. 47, 2005.

[20] T. Bailey, J. Nieto, J. Guivant, M. Stevens, and E. Nebot, "Consistency of the EKF-SLAM algorithm," 2006, pp. 3562-3568.

[21] A. Barrera and A. Weitzenfeld, "Biologically-inspired robot spatial cognition based on rat neurophysiological studies," *Autonomous Robots,* vol. 25, pp. 147-169, 2008.

[22] D. Ball, S. Heath, G. Wyeth, and J. Wiles, "iRat: Intelligent rat animat technology," in *Australasian Conference on Robotics and Automation 2010*, G. Wyeth and B. Upcroft, Eds., ed Brisbane, Australia, 2010.

[23] M. Quigley, B. Gerkey, K. Conley, J. Faust, T. Foote, J. Leibs, E. Berger, R. Wheeler, and A. Ng, "ROS: an open-source Robot Operating System," in *Proceedings of the Open-Source Software workshop at the International Conference on Robotics and Automation (ICRA)*, ed, 2009.

[24] A. Censei, "An ICP varient using a point-to-line metric," presented at the The IEEE International Conference on Robotics and Automation (ICRA), Pasedena, CA, 2008.

[25] A. Cangelosi, A. Greco, and S. Harnad, "From robotic toil to symbolic theft: grounding transfer from entry-level to higher-level categories1," *Connection Science,* vol. 12, pp. 143-162, 2000.

[26] T. Nagel, "What Is It Like to Be a Bat?," *The Philosophical Review,* vol. 83, pp. 435-450, 1974.

[27] A. C. Varzi, "Spatial reasoning and ontology: parts, wholes, and locations," *Handbook of Spatial Logics,* pp. 945-1038, 2007.