

THE FORMATION, GENERATIVE POWER, AND EVOLUTION OF TOPONYMS: GROUNDING A VOCABULARY IN A COGNITIVE MAP

RUTH SCHULZ, DAVID PRASSER, PAUL STOCKWELL, GORDON WYETH, AND
JANET WILES

*School of Information Technology and Electrical Engineering,
The University of Queensland,
Brisbane, QLD, 4072, Australia*

We present a series of studies investigating the formation, generative power, and evolution of toponyms (i.e. topographic names). The domain chosen for this project is the spatial concepts related to movement through the environment, one of the key sets of concepts to be grounded in autonomous agents. Concepts for spatial locations cannot be directly perceived and require representations built from interactions and inferred from ambiguous sensory data. A generative toponymic language game has been developed to allow the agents to interact, forming concepts for locations and spatial relations. The studies have shown that a grounded generative toponymic language may form and evolve in a population of agents interacting through language games. Initially, terms are grounded in simple spatial concepts directly experienced by the robots. The generative process then enables the robots to learn about and refer to locations beyond their direct experience, enabling concepts and toponyms to co-evolve.

1. Introduction

The challenge of understanding the evolution of language is not just in understanding the “frozen accident” of how it did evolve in humans, but how it could evolve. In order to investigate this question, one needs to understand not just what *is* but what in principle *could be*. Autonomous agents provide a methodology for investigating such issues where mechanisms designed as engineering solutions can inform universal principles.

The relationship between concepts and how they are grounded in embodied language is one of the key questions in the philosophy of language, and is central to our use of language in embodied agents. The domain chosen for study in this project is the spatial concepts related to movement through the environment, one of the important concepts for mobile autonomous agents. The position taken for this project is that language must be grounded to be used meaningfully; language

is more useful if it is generative; and concepts and language are formed interactively.

Grounding can be described as the link between things in the world and internal categories (Cangelosi, 2006), and how words and speech are related to the language user's environment (Roy, 2005). The importance of grounding for language is to provide meaning for primary concepts and the ability to associate language terms with those concepts. Human language is generative rather than being a one-to-one labelling of symbols to concepts. An agent with a generative language can refer to concepts not yet experienced, and ground concepts that cannot be directly experienced by extension from existing concepts.

The most basic spatial concepts correspond to areas in space and are referred to by labels for places, such as city or suburb names. Areas within an environment or along a path can also often be described by single words, such as corner, corridor, or intersection, or larger regions such as kitchen, office, or backyard. We call names for specific places in an environment *toponyms* (i.e. topographic names), and a set of such terms to comprehensively describe an agent's environment we term a *toponymic language*.

When people describe spatial locations, landmarks are preferred, followed by spatial relations (Tversky, 2003). In English, spatial relations are generally provided by spatial prepositions, with directions and distances combined to form spatial terms such as 'in front of', 'near', and 'at'. Other languages and cultures express spatial relations in various ways (Levinson, 1996), for example, the Mayan language Tzeltal has only one preposition with nouns and verbs providing spatial location information (Brown, 2006). Human experiments (Logan & Sadler, 1996) and theoretical investigations (O'Keefe, 1996; Zwarts, 1997) have described spatial templates defining areas in the world, relative to references, where spatial terms are used. Several models of spatial language have been developed, including two studies that involve terms related to spatial locations (Bodik & Takac, 2003; Steels, 1995). The language games of these studies involve concepts of direction and distance from the agent to an object in the world. In these studies, the locations of the objects and agents are unambiguous and known by both agents.

There is a natural human tendency to assume that a spatial language entails descriptions of *objects* at specific locations, or using objects to define landmarks. From autonomous robot studies it is clear that a set of concepts to describe space and a corresponding toponymic language does not require knowledge of objects, nor descriptions of visual scenes (Milford, Schulz, Prasser, Wyeth, & Wiles, 2007). Using an autonomous agent methodology has the benefit of making it possible to investigate the grounding of spatial languages without objects.

The challenge for this project is to combine grounding and generative languages. This may be achieved by forming a generative language in embodied agents with representations obtained from experiences with the world. The representations required to form spatial concepts from movement through the world can be formed from an established robot platform. As spatial locations cannot be directly perceived, the representations must abstract from direct sensory inputs to allow knowledge about locations relative to other locations in the world. RatSLAM (Milford et al., 2007) is a robotic platform that meets these requirements. Each agent has a unique representation of the world based on their own experiences. The challenge is for two or more agents, each with unique representations, to learn to communicate with each other. Language games can be played to form concepts from these representations through interactions with the world and other agents by naming spatial concepts while exploring the world.

The overall goal of the project is to explore issues in the relationship between language, concepts, and grounding in autonomous agents with respect to spatial locations. The specific aims are to show that mobile robots can form toponymic concepts, that these concepts can be formed indirectly through a generative language, and can be learned and used by successive generations.

Three studies are presented in this paper that investigated the formation, generative power, and evolution of toponyms. In the first study robots in a simulation world played a toponymic language game. In the second study agents played a generative toponymic language game in a grid world. The third study investigated the evolution of a toponymic language in a grid world with generations of agents that played language games.

2. Study 1. Formation of Toponyms

The basic spatial concepts of areas in space can be used to ground other spatial concepts. Shared attention for these concepts is being located in the ‘same place’. In previous studies, language games have been used to name objects (Bodík & Takáč, 2003) and other agents (Steels, 1995) in the world, but not to form concepts and names for areas that are undefined prior to the interactions of the agents. In this first study the formation of toponyms and scaling effects are investigated in a simulation world (see Figure 1) with two agents. In toponymic language games, agents play a game whenever they are within hearing distance of each other. One agent is the speaker and the other agent is the hearer. The speaker agent chooses the best word for its current location, and the hearer agent updates its lexicon with this word and its current location. Three conditions were tested, based on hearing distances of 1m, 3m, and 5m.




Figure 1 Simulation world map and robot view. The robot world is an open plan office. In the map, the black hexagons are desks.

The representations of the world are nodes or experiences in an experience map unique to each robot that was constructed during an exploration phase. The experience map is an approximate x - y representation of the world that each robot learns from its visual information and odometry. At any point in time one experience in the map is active, encoding the robot's best estimate of its position. The robot mapping model used, RatSLAM, is based on the rodent hippocampus. (For more information, see Milford et al., 2007).

A lexicon table associates the experiences of the robot and distinct words. The association between an experience and a word is increased by *incr* every time they are used together in a game. A most information strategy is used to choose a word for a location. In this strategy, the agent uses the association between each experience and word, A , and the sum of how associated each word is with all experiences, C . For the chosen location, the information transmitted by each word, I , is calculated by dividing the sum of A for all n nearby experiences by C (see Eq. 1), where experiences are nearby the current experience if they are within x meters in the experience map of the robot. The distance x is set to the hearing distance used, with distances in the experience map being approximately equivalent to distances in the simulation world. The word with the highest value of I is chosen. When all of the experiences associated with a word are nearby the location, I is 1.0. If I is less than *thresh*, a new word is invented. In this study, *incr* is 1.0 and *thresh* is 0.2.

$$I_i = \frac{\sum_{j=1}^n A_{ij}}{C_i} \quad (1)$$

2.1. Results and Discussion

In all three conditions, the robots developed a shared set of toponyms (see Figure 2), showing that toponyms can be formed at different levels of scale by using a toponymic language game with different scales of interaction between agents. In the case where the toponyms cover larger areas, the agents can choose a different word for the same location in the world, even when the templates of

the toponyms are similar between the agents. At a larger scale, the exact location of the toponym and the borders between toponyms are less certain.




Figure 2 Toponym templates, shown relative to the experience map of one of the agents, for hearing distances of a) 1m, b) 3m, and c) 5m. The hearing distance of the robot affected the number of words invented, as with a larger hearing distance, toponyms cover larger areas.

3. Study 2. Generative Power of Toponyms

From the first study, we know that toponyms can be formed for all places in the world visited by both agents, by playing toponymic language games when within hearing distance of each other. The next step is forming relations between these toponyms, and using these relations to generate concepts and labels for places that cannot or have not been visited by the agents. The second study investigated the formation of toponyms and spatial relations with two agents that played generative toponymic language games. In these games, words are formed for the current location, an orientation location, a target location, and the spatial words of direction and distance describing the target location (see Figure 3a). The agent is at the current location facing the orientation location and talking about the target location by specifying the direction and distance to the target.

Each location in the grid is equivalent to an experience used in Study 1. Distance and direction are calculated from the current, orientation, and target locations. The world is a grid world, used for computational efficiency, in which agents occupy any location without an obstacle (see Figure 3b,c). The hearing distance for the agents is the four nearest neighbour locations. Two conditions are tested based on the empty world and the world with desks. The study consisted of five runs of 20000 games for each condition.




Figure 3 a) The elements involved in a generative language game: The agent is at 'Current' facing 'Orientation' and talking about 'Target'; toponyms are found for the current, orientation, and target locations, and spatial words are found for the direction, θ , and distance, d . b) Empty grid world map of size 15x15 c) Grid world map of size 15x15 with desks similar to the world of the previous study

Toponyms are comprehended by forming the normalised template for the toponym where the best location has a value of 1.0, and a location with no association with the word has a value of 0.0. The success, s_T , of the toponym for the current location is the value of the template at the current location. Toponyms are produced as in the previous study. Instead of using a threshold to control the invention of new words, this now occurs according to the success of how well the speaker comprehends its own utterance, with a probability, p , as shown in Eq. 2, with $T=0.3$. This value of T allows the agents to invent words when the success is low and have a stable lexicon at higher levels of success.

$$p = \exp\left(\frac{-1}{(1-s_T)T}\right) \quad (2)$$

The direction and distance lexicons of the agents each have 50 values that words may be associated with, corresponding to 50 ranges of directions and distances. When producing or comprehending spatial words, the agent may choose the level of generalisation used in the size of the neighbourhood of values considered. When the target location is specific, a small neighbourhood is used (e.g. 1-5 values), and when the target location is more general, a large neighbourhood is used (e.g. over 30 values). Spatial words are comprehended by forming the templates for the target toponym, t , and for the spatial words, t_s . The success of the generative game, s_G , is found by comparing how well these templates match each other over all locations in the world, as shown in Eq. 3.

$$s_G = \sum_{i=1}^{15 \times 15} \min(t_{ti}, t_{si}) \quad (3)$$

Every time a game is played, the lexicon tables of the hearer are updated. The speaker's lexicon is updated when a new word is invented. The toponym lexicon table is updated as in the previous study for the current location. The templates of the target location and the spatial words are used to update the lexicon tables for the target toponym and spatial words, adding a total of $incr$ to the lexicon associations across the locations or values. In this study, $incr = 1.0$.

3.1. Results and Discussion

In both the empty world and the world with desks, the rate of word invention was highest for the first 100 games, and agents continued to invent words throughout each run. The average final lexicon in the empty world included 24.2 toponyms, and in the world with desks included 34.8 toponyms. The average success of the

agents in their final 100 games over the five runs was 0.71 for toponyms and 0.63 for generative toponyms for the empty world, and 0.76 for toponyms and 0.64 for generative toponyms for the world with desks.

The toponyms invented and used by the agents in the empty world are all specific (see Figure 4). Most of the toponyms for the agents in the world with desks are also specific with a few exceptions that are associated with larger areas of the world. The distance and direction words of the agents in both worlds are similar, although the way in which they are used varies. Over the five runs, 58.2% of games in the empty world and 42.7% of games in the world with desks used a low generalization (1-5 values); and 14.3% of games in the empty world and 42.8% of games in the world with desks used a high generalization (>30 values).




Figure 4 Templates for a) toponyms, b) distances, and c) directions for an agent in the empty world, showing the lexicon with the neighborhood function applied. The size of the neighborhood function used in the figure is the average size used by that agent over the run.

This study has shown that agents can form concepts for spatial relations describing the relative locations of toponyms, and use these to form new toponyms. The new toponyms tended to be general and resulted in the use of the direction and distance terms in a more general way.

4. Study 3. Evolution of Toponyms

From the second study, we know that generative toponymic languages with concepts for places in the world and relations between these places can be formed for two agents playing toponymic language games. The question now is how this language might change over generations of agents. The third study investigated the evolution of a generative toponymic language.

The words, concepts, methods of production and comprehension, and measures of success were the same as in the previous study. The world was the 15 by 15 grid with desks. In the initial population there were two agents as in the previous study. In subsequent generations, the older agent was removed from the population and a new agent was added, initially as a hearer. After 500 games, the

new agent played games as either a speaker or a hearer. In this study, two conditions are tested based on generations of 1000 and 2000 games. Both conditions consisted of five simulation runs of 20000 games.

4.1. Results and Discussion

The first generation for each run formed their language through negotiation, in which the success of the toponymic and generative games slowly increased as the languages were formed (see Figure 5). New agents in later generations began by learning from the older agent, which caused a drop in success that quickly returned to a high level as the new agents learned the language. Most of the toponyms are specific toponyms. Over generations of agents, specific toponyms tended to remain stable, as did the concepts for directions and distances while the more general toponyms shifted their meaning to become more specific (see Figure 6). Agents may evolve a more specific toponymic language for locations that cannot be visited if allowed to form the language over generations of agents.




Figure 5 Average success for Toponym and Generative Language Games over 5 runs calculated every 100 games for a) 1000 game generations and b) 2000 game generations




Figure 6 Four toponyms over 20 generations of 1000 games. Each square from left to right is the template for this word for the agent that left the population at this generation. Specific toponyms do not change much over time while toponyms that initially refer to two distinct locations will come to refer to only one of these. Some become associated with more general areas, due to the uncertainty involved in updating the lexicon. These concepts may shift over time.

5. General Discussion and Conclusion

The studies in this paper have shown how a generative toponymic language may form and evolve in a population of agents. Agents were able to form concepts for locations, directions, and distances as they interacted with each other and associated words with underlying values. The relations between existing concepts were used to expand the concept space to new locations. Evolution allowed the general toponyms referring to new locations to become more specific. A generative language can be formed with concepts for which the dimensions for extension are clear, in this case in an approximate x - y representation of the world.

These studies have combined grounding, generative language and spatial locations. Spatial locations were grounded in a robot's experiences obtained from exploring the world. A generative toponymic language game was introduced that enables the grounding of concepts of locations beyond direct experience. The parallel formation of concepts and words allowed concepts and words to have different levels of generalisation, and showed how the concepts can be formed through social interactions.

A functional perspective of getting robots to talk to each other about spatial locations was used to inform about grounding and generative languages. We are currently extending this study into the simulation world of the robots, and investigating other concepts, including verbs describing the robots motion through the world.

Acknowledgements

RS and PS were supported by Australian Postgraduate Awards. This research is funded in part by a grant from the Australian Research Council.

References

- Bodík, P., & Takáč, M. (2003). Formation of a common spatial lexicon and its change in a community of moving agents. In *Frontiers in AI: Proceedings of SCAI'03*: IOS Press.
- Brown, P. (2006). A sketch of the grammar of space in Tzeltal. In S. C. Levinson & D. Wilkins (Eds.), *Grammars of Space: Explorations in Cognitive Diversity* (pp. 230-272). Cambridge: Cambridge University Press.
- Cangelosi, A. (2006). The grounding and sharing of symbols. *Pragmatics and Cognition*, 14(2), 275-285.
- Levinson, S. C. (1996). Language and Space. *Annual Review of Anthropology*, 25, 353-382.

- Logan, G. D., & Sadler, D. D. (1996). A computational analysis of the apprehension of spatial relations. In P. Bloom, M. A. Peterson, L. Nadel & M. F. Garrett (Eds.), *Language and Space*. Cambridge, Massachusetts: The MIT Press.
- Milford, M., Schulz, R., Prasser, D., Wyeth, G., & Wiles, J. (2007). Learning spatial concepts from RatSLAM representations. *Robotics and Autonomous Systems - From Sensors to Human Spatial Concepts*, 55(5), 403-410.
- O'Keefe, J. (1996). The spatial prepositions in English, vector grammar, and the cognitive map theory. In P. Bloom, M. A. Peterson, L. Nadel & M. F. Garrett (Eds.), *Language and Space*. Cambridge, Massachusetts: The MIT Press.
- Roy, D. (2005). Semiotic Schemas: A framework for grounding language in action and perception. *Artificial Intelligence*, 167(1-2), 170-205.
- Steels, L. (1995). A self-organizing spatial vocabulary. *Artificial Life*, 2(3), 319-332.
- Tversky, B. (2003). Places: Points, planes, paths, and portions. In E. van der Zee & J. Slack (Eds.), *Representing direction in language and space* (pp. 132-143). New York: Oxford University Press, Inc.
- Zwarts, J. (1997). Vectors as relative positions: a compositional semantics of modified PPs. *Journal of Semantics*, 14, 57-86.