THE FORMATION, GENERATIVE POWER, AND EVOLUTION OF TOPONYMS: GROUNDING A SPATIAL 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 places in an environment, one of the key sets of concepts to be grounded in autonomous agents. Concepts for places cannot be directly perceived as they require knowledge of relationships between locations in space, with representations inferred from ambiguous sensory data acquired through exploration. A generative toponymic language game has been developed to allow the agents to interact, forming concepts for locations and spatial relations. The studies demonstrate how a grounded generative toponymic language can form and evolve in a population of agents interacting through language games. Initially, terms are grounded in simple spatial concepts directly experienced by the agents. A generative process then enables the agents to learn about and refer to locations beyond their direct experience, enabling concepts and toponyms to co-evolve. The significance of this research is the demonstration of grounding for both experienced and novel concepts, using a generative process, applied to spatial locations.

1. Introduction

For autonomous agents to interact effectively with humans, they require the ability to connect their internal representations of the world to human language. Grounding refers to the relationship between things in the world, internal categories, and their symbols (Harnad, 1990). While researchers have emphasised different aspects of the grounding problem, the central role of grounding is to provide meaning for primary concepts and to associate language terms with those concepts. Our approach emphasises interaction between concepts and language, rather than the primacy of one or the other. Human language is generative rather than being a one-to-one labelling of symbols to concepts. Hence a complete theory requires the grounding of concepts that cannot be directly experienced. Appropriate representations are a way of bridging between symbols and the world. In particular, a cognitive map provides an internal representation of places and their relations in the world (O'Keefe & Nadel, 1978).

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 or corridor, or larger regions such as kitchen or office. We call names for places in an environment *toponyms* (i.e. topographic names), and a set of such terms to comprehensively describe an environment a *toponymic language*.

In this study, we have drawn on insights from behavioural studies of spatial language, related mathematical and computational models, and agent-based language games. In English, spatial relations are generally referred to by spatial prepositions, with directions and distances combined to form spatial terms such as ‘in front of’, ‘near’, and ‘at’. Human experiments (Logan & Sadler, 1996) and theoretical investigations (O'Keefe, 1996; Zwarts, 1997) have described spatial templates for terms defining areas in the world. Models of spatial language have been developed, including language game studies where agents formed a vocabulary for predefined concepts of agents and spatial relations (Steels, 1995), and where a shared spatial language emerged to describe directions, distances, and object names (Bodik & Takac, 2003). Studies to date that have demonstrated grounding in a spatial domain have used location concepts that were unambiguous and known by all agents, and an absolute direction system, where all agents know the reference direction.

The challenge for this project is to combine grounding and generative languages by forming a generative language in embodied agents. 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, Schulz, Prasser, Wyeth, & Wiles, 2007) is a robotic platform that meets these requirements. The objective is for two or more agents, each with unique representations of the world based on their own experiences, to learn to communicate with each other, and to be able to direct each other to locations. Language games can be played to form concepts from these representations through interactions with the world and other agents.

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 autonomous agents can form toponymic concepts and vocabulary, that both concepts and labels can be formed indirectly through a generative process, and can be learned and used by successive generations.

Three studies were designed to investigate the formation, generative power, and evolution of toponyms. In the first study, autonomous agents (simulated robots) played a toponymic language game. In the second study, the toponymic language game was extended to include a generative task. The third study investigated the evolution of the language over generations.

2. Study 1. Formation of Toponyms

The basic spatial concepts of areas in space require an understanding of locations. For the first study, we designed a spatial naming game to investigate the formation of toponyms and scaling effects in a simulation world (see Figure 1a,b) with two agents. In toponymic language games, agents interact whenever they are within hearing distance of each other. The speaker agent chooses the best word for its current location, and the hearer agent updates its lexicon.

In the RatSLAM system, each robot learns a unique representation of the world as a topological map of experiences, constructed during an exploration phase (see Figure 1). An experience map is an approximate *x-y* representation of the world that each robot constructs 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 (for more information, see Milford et al., 2007).

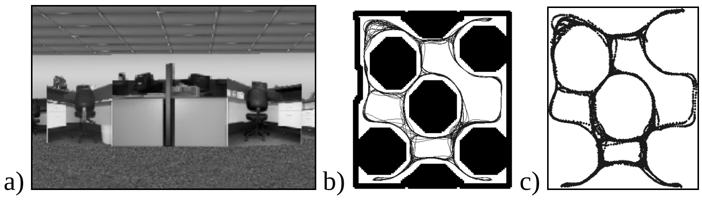

Figure 1 a) Simulation robot view, b) World map, and c) Experience map. The world is an open plan office. In the map, the black hexagons are desks, and the path of the robot is shown. In the experience map, each dot shows the location of an experience in the robot's internal map.

A lexicon table stores the associations between the experiences of the robot and distinct words. The association between an experience and a word is strengthened when they are used together. For each location the word with the highest information value is chosen. The information value, I_{wp}, for the word, w, in location, p, is the relative information of the word within a neighbourhood of size D compared to the total usage of the word, calculated as follows:

$$I_{wp} = \left(\sum_{n=1}^{N} A_{wn}(D-d_{np})/D\right) \Big/ \sum_{m=1}^{M} A_{wm} \qquad (1)$$

where N is the number of experiences within D of the location, p; A_{wn} is the association between the word, w, and an experience, n; d_{np} is the distance between an experience, n, and the location, p; and M is the total number of experiences in the robot's experience map.

In each interaction, words are invented with probability, p, as follows:

$$p = e^{-1/(1-S)T} \tag{2}$$

where S is the success of the interaction, equal to the information value of the location-word combination, and T is the temperature, which sets the level of success accepted by an agent. Using a word invention rate corresponding to the success of the interaction allows agents to use words where they provide significant information about the current location, and to invent words otherwise. Varying the temperature alters the rate of word invention, where a higher temperature increases the probability of inventing a new word.

Our study used simulated agents rather than robots with a hearing distance of 3m, and a neighbourhood size, D, of 5m. Within a trial, the temperature for word invention was set at a fixed temperature, T, and agents evolved a set of words. Three conditions were tested, based on low, medium, and high temperatures, with each condition run for 2000 interactions. In all three conditions, the agents developed a shared set of toponyms (see Figure 2), showing that toponyms can be formed at different levels of scale by using different rates of word invention. Each location is referred to by a toponym in its vocabulary, interpreted as the most informative point on the experience map. A higher temperature resulted in a more specific toponymic language. The study demonstrated how toponyms could be formed for all places in the world visited by both agents, by playing toponymic language games when within hearing distance.

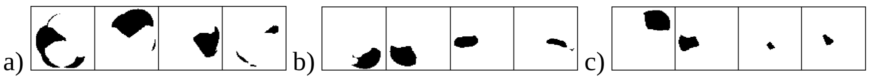

Figure 2 Toponym meanings shown as toponym usage templates. Each set (a-c) shows four of the words for one agent from a trial. Each cell shows the locations in the experience map of the agent where the word is used. (a) The lowest temperature, T=0.25, resulted in the smallest number of words, with four of the five words covering large areas; (b) The medium temperature, T=0.5, resulted in 18 words, with 11 covering large or medium areas; (c) the highest temperature, T=0.75, resulted in the greatest number of words, 28, with 21 covering small areas.

3. Study 2. Generative Power of Toponyms

To go beyond simple concepts requires a generative process. In the second study relations are formed between toponyms, and used to generate concepts and labels for places that cannot or have not been visited by the agents. A key challenge for embodied language games is to take into account the different perspectives of the agents. The generative toponymic language game, adapted from previous language games (Bodik & Takac, 2003; Steels, 1995), is based on naming three locations: Both agents are located within hearing distance at the first (current) location, they are facing the second (orientation) location, hence aligning their perspectives, and then they talk about a third (target) location (see Figure 3a). Given the three locations, agents can describe the target location with spatial words of distance and direction.

For computational tractability, the second study used a simple grid world (see Figure 3b,c). Each agent's experience map is simulated by a corresponding grid of experiences, with each location in the grid equivalent to an experience used in Study 1.

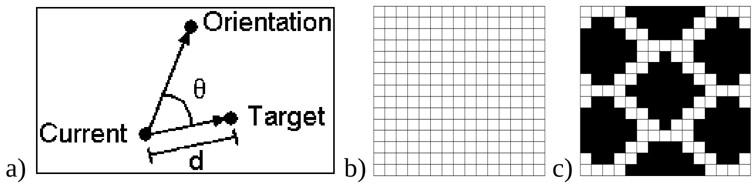

Figure 3 a) The elements involved in a generative language game: The agent is at Current facing Orientation and talking about Target; toponyms are selected for the current, orientation, and target locations, and spatial words are selected for the direction, θ, and distance, d. b) Empty grid world map of size 15x15 c) Grid world map of size 15x15 with desks similar to the world of Study 1.

Each toponym has a corresponding template, which is calculated from the association between the toponym and all nodes in the experience map. The experience with the strongest association has a value of 1.0. The success of the toponym for an interaction is the value of the toponym template for the experience being used by the agent for the interaction. Toponyms are selected and invented as in Study 1, with the neighbourhood for calculating information being the four nearest neighbour locations. The probability of inventing new words is calculated as in Study 1.

The direction and distance lexicon tables of the agents are vectors of 50 values that words are associated with, corresponding to a range of directions and distances. Each combination of the spatial words of distance and direction words has a corresponding template which is calculated from the associations between the spatial words and the vectors of values. The spatial words forming the template that best matches the target toponym template are selected by the speaker. The success of the generative interaction is calculated by comparing the templates for the target toponym and the spatial words. The probability of inventing spatial words is calculated as for the toponyms using the success of the generative interaction.

Every time the agents interact, the lexicon tables of the hearer are updated. The speaker's lexicon is updated when a new word is invented. The templates of the target location and the spatial words are used to update the lexicon tables for the target toponym and spatial words, increasing the lexicon associations across the experiences and vectors of values.

In this study, two conditions were tested based on the empty world and the world with desks. The hearing distance for the agents was the four nearest neighbour locations. The temperature, T, was 0.25, which allowed a level of specificity for toponyms of 5-10 experiences. The study consisted of five trials of 10,000 interactions for each condition.

In both the empty world and the world with desks, the rate of word invention was highest for the first 100 interactions, and agents continued to invent words throughout each trial. The toponyms invented and used by the agents in the empty world were all specific, and some of the toponyms used by agents in the world with desks were general (see Figure 4). The average final lexicon in the empty world had 27.8 toponyms, and in the world with desks had 31.4 toponyms. There were more toponyms in the world with desks because they include the general toponyms, which cover similar areas.

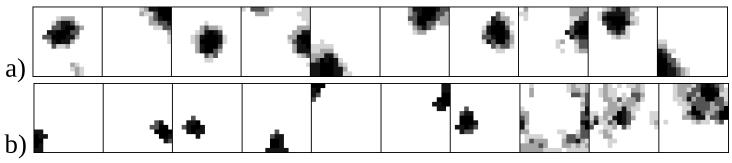

Figure 4 Toponym templates. Non-white regions show that the word is one of the top five words providing information about a location, with black indicating that the word will be used at a location. Each set (a-b) shows templates for 10 of the words for one agent from a trial. a) In the empty world all templates were specific; b) In the world with desks, most templates were specific, but some were general, formed by referring to a location through the generative process.

4. Study 3. Evolution of Toponyms

Languages are not just created within a single agent's lifetime. They evolve and are refined over generations of agents. The third study investigated the evolution of a generative toponymic language.

The words, concepts, selection of words, comprehension, and measures of success were the same as in the Study 2. The world was a 15 by 15 grid with desks (see Figure 3c). Generations consisted of a set number of interactions, g. In the initial population two agents play negotiation games. In subsequent generations, the older agent was replaced by a new agent, initially as a hearer. After *g*/2 interactions, the new agent could interact as a speaker or a hearer. In this study, two conditions were tested based on g = 1000, and g = 2000, each consisting of five trials of 20,000 interactions.

The first generation for each trial formed their language through negotiation, in which the success of the toponymic and generative games increased as the languages were formed (see Figure 5e). Over generations, specific toponyms tended to remain stable, as did the concepts for directions and distances while the more general toponyms shifted to become more specific (see Figure 5a-d). The results presented are for the first ten generations of the condition where g = 1000. Similar results were obtained for the remainder of the generations and for the condition where g = 2000.

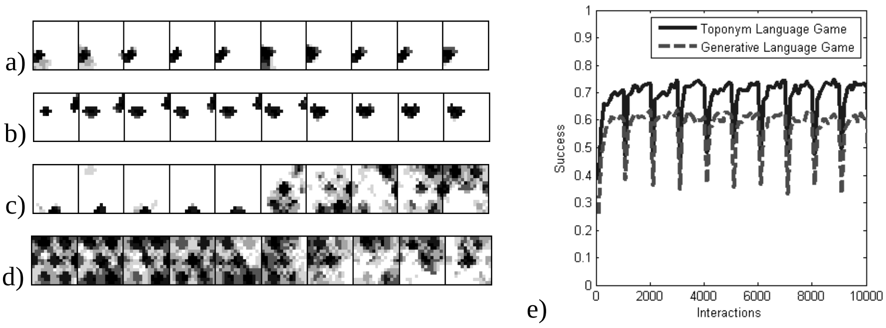

Figure 5 Language games over generations. (a-d) Toponym templates over generations. Each row shows how a toponym is used throughout the trial, with each cell being the toponym's template for the agent leaving the population at that generation. Each row (a-d) is an example of different types of toponyms: a) shows a specific toponym that does not alter much throughout the generations; b) shows a toponym that initially refers to multiple specific locations, but only refers to one of these after several generations; c) shows a specific word that becomes more general; d) shows a general word that became more specific. e) Success of language games over generations. The success of a toponym language game is the information value of the word used for the current location. The success of a generative language game is how well the toponym template matches the spatial words template for the words used. The peak average success was just over 0.6 for the generative language game, and just over 0.7 for the toponym language game. As a new agent entered the population, they 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.

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. 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.

The key contribution of the research is the demonstration of grounding for both experienced and novel concepts using a generative process, applied to spatial locations. We have shown that generative grounding can be achieved with an appropriate representation of the concept space (in this case, an approximate *x*-*y* representation of the world), a way to form and label intrinsic concepts (in this case, toponyms), and a generative process that creates both the concepts and the labels. We are currently extending this study into the simulation world, and investigating other concepts, including verbs describing the robot's 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

Bodik, P., & Takac, M. (2003). Formation of a common spatial lexicon and its change in a community of moving agents. In B. Tessem, P. Ala-Siuru, P. Doherty & B. Mayoh (Eds.), *Eighth Scandinavian Conference on AI*. Amsterdam: IOS Press.

Harnad, S. (1990). The symbol grounding problem. *Physica D: Nonlinear Phenomena, 42*, 335-346.

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.

O'Keefe, J., & Nadel, L. (1978). *The hippocampus as a cognitive map*. New York: Oxford University Press.

Steels, L. (1995). A self-organizing spatial vocabulary. *Artificial Life, 2*(3), 319-332.

Zwarts, J. (1997). Vectors as relative positions: a compositional semantics of modified PPs. *Journal of Semantics, 14*, 57-86.