Language Change across Generations for Robots using Cognitive Maps

Ruth Schulz[1], Gordon Wyeth[2] and Janet Wiles[1]

[1] School of Information Technology and Electrical Engineering,
The University of Queensland
[2] Queensland University of Technology
ruth@itee.uq.edu.au

Abstract

Languages change over time, as new words are invented, old words are lost through disuse, and the meanings of existing words are altered. The processes behind language change include the culture of language acquisition and the mechanisms used for language learning. We examine the effects of language acquisition and learning, in particular the length of the learning period over generations of robots. The robots form spatial concepts related to places in an environment: toponyms (place names) and simple prepositions (distances and directions). The use of spatial concepts allows us to investigate different classes of words within a single domain that provides a clear method for evaluating word use between agents. The individual words used by the agents can change rapidly through the generations depending on the learning period of the language learners. When the learning period is sufficiently long that more words are retained than invented, the lexicon becomes more stable and successful. This research demonstrates that the rate of language change depends on learning periods and concept formation, and that the language transmission bottleneck reduces the retention of words that are part of large lexicons more than words that are part of small lexicons.

Introduction

Language change is a ubiquitous property of natural languages. One characteristic of language change is the production of neologisms, with new words created or existing words modified, combined, or separated (Brinton & Traugott, 2005). A shared language can be sustained within generations, while the words and concepts may change through generations. Although older generations are prone to deplore the language of younger generations, language change only becomes a problem when members of a population are no longer able to understand each other (Aitchison, 1991).

There are three timescales on which language change occurs: individual learning, cultural transmission, and biological evolution (Kirby, Dowman, & Griffiths, 2007). Language change is driven by both external sociolinguistic and internal psycholinguistic factors (Aitchison, 1991). Constraints that shape language include sensorimotor factors (the noisiness and variability of signals), cognitive limitations (learning, processing, and memory), thought (concepts and categorization), and pragmatic constraints (Chater & Christiansen, 2009). Language acquisition mechanisms influence the nature of language change (Niyogi, 2006), with the transmission of language from one generation to the next involving the mechanisms of language learning and production (Brighton, Smith, & Kirby, 2005).

Representation and culture influence the concepts that can be formed in a language and the ease with which agents learn these concepts. These factors are part of concept formation, language production, and language acquisition mechanisms. Together with learning mechanisms, representation affects how individual agents form concepts, which in turn affects the concepts that form in a population of agents. The cultural environment of the agents determines the words and concepts that agents are exposed to over their lifetimes.

A variety of representations and learning mechanisms have been used in studies investigating language evolution in computational agents. Recent studies have investigated the use of visual perceptions and spatial representations in forming a language for regions in geographical space and generative grounding using spatial representations (Schulz, Prasser, Stockwell, Wyeth, & Wiles, 2008). When agents ground concepts generatively, by combining existing concepts to form new concepts, there is increased flexibility and hence also ambiguity in the association between words and concepts.

In generational studies, agents start afresh with each new generation, learning the existing language and potentially expanding it. A reason that language is evolvable is that it is situated in a cultural environment that aids learning through generations, which can be implemented with iterated learning (Brighton, et al., 2005; Kirby & Hurford, 2002), in which agents learn language from the utterances of other agents. The strategies used by language speakers and hearers in determining what to talk about and how to talk about it are also a part of culture.

One feature of culture that has been studied previously is the bottleneck of language transmission (Brighton, et al., 2005; Kirby, 2002; Smith, 2007; Tonkes & Wiles, 2002). The bottleneck has been found to be important for the development of compositional and productive language. Previous spatial language studies have investigated how the rate at which agents enter and leave the population affected whether the agents were able to sustain a shared spatial language (Bodík & Takáč, 2003). These studies found that when the length of time agents spent in the population was sufficiently long (i.e. the bottleneck was sufficiently large), a shared spatial language was able to be sustained. These results have also been found in language studies with arbitrary feature representations (Smith, 2007).

Studies investigating the language transmission bottleneck have either considered a single class of words or analyzed the success of the whole language, with individual words used as examples. However, different classes of words, such as nouns and prepositions, play different roles in meaningful communication, and all classes of words may not be equally likely to pass through a language bottleneck.

The challenge for this project is to determine how spatial languages can change through generations and to determine how the length of the learning period and lifetime of the agents affect language change. The main questions to answer include how to interpret spatial language change over generations and whether different types of spatial words have different rates of change. In particular we are interested in how learning by successive generations affects the turnover of individual words. The study described in this paper investigated the effect of the length of the learning period and the lifetime of the agents on the various spatial concepts that form and how the language changes throughout the generations.

A Spatial Language with Cognitive Maps

In language studies, the agent interactions influence the words and concepts that a language agent is exposed to and chooses to use throughout its lifetime. The specific games played determine which niches of concept space will be filled and the words chosen by the agents determine which words will survive through generations. In the study presented here, generations of simulated robots played language games to form concepts for toponyms (place names) and simple prepositions (directions and distances). The length of each generation was varied from four interactions per generation up to 1000 interactions per generation to investigate the affect of the length of the learning period and agent lifetimes on language change. The nature of the language change was investigated by comparing rates of word invention, retention, and persistence for the different concept types of toponyms, directions, and distances.

Location Language Games

The language games used in these studies are location language games (see Figure 1). To play a location language game, the agents require a representation of the world acquired through exploration carried out independently of other agents in the world. Shared attention for location language games is co-presence, that is, the agents are within hearing distance. While autonomously exploring the world, the agents intermittently send a "Hello" signal. If a "Hello" signal is heard, the hearing agent sends a "Hear" signal and the agents play a game. After shared attention is established, the speaker chooses a topic, which in a location language game relates to the current location of the agents or a location at a distance from the agents, depending on the game being played. After the topic is determined, the speaker uses its lexicon to determine which word should be used in the current situation and produces an utterance. Both agents then update their representations and lexicon. In the location language

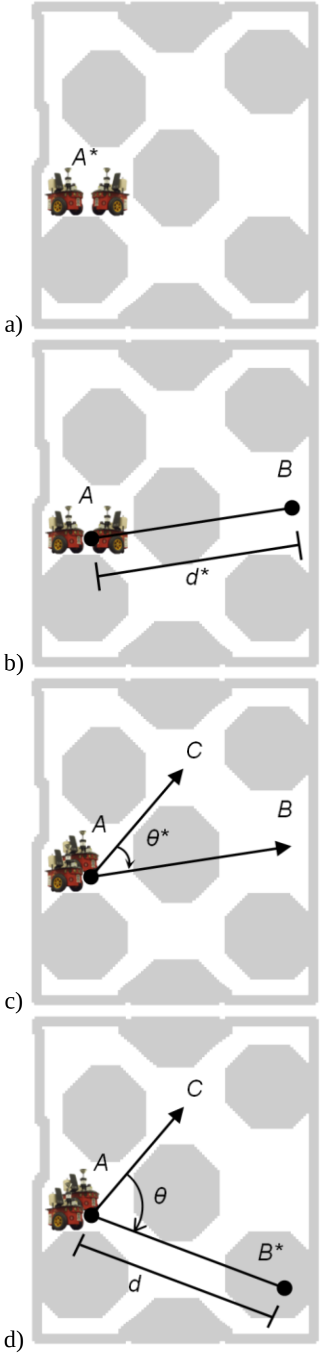

Figure 1. Referents used in the language games: a) The *where-are-we* game involves a single location: the current location, *A*, of both robots. b) The *how-far* game involves two locations (current, *A*, and target, *B*) and a distance, *d*. c) The *what-direction* game involves three locations (current, *A*, target, *B*, and orientation, *C*) and a direction, *θ*. d) The *where-is-there* game involves three locations (current, *A*, target, *B*, and orientation, *C*), a direction, *θ*, and a distance, *d*. The figures show the robots located in the open plan office of the simulation world, with gray lines representing walls and gray octagons representing desks. A star (*) indicates that the speaker may invent a new word and that both agents will update their lexicon for the marked word.

games played in this study the hearer receives the utterance and updates their representations, but does not explicitly evaluate the speaker's utterance and no feedback is given to either agent. Repeated encounters enable coherent languages to form even without explicit feedback (a phenomenon reported in a variety of studies including Smith, 2007; and Vogt, 2004).

In the study, the agents played *where-are-we*, *how-far*, *what-direction*, and *where-is-there* games. The premise of a *where-are-we* game is a location language game where the topic is the current location of the agents (see Figure 1a). The speaker produces a word for the current location and both agents update their lexicon based on the speaker's utterance.

The *how-far* game is based on naming two locations: Both agents are located at the first location (A) and they talk about the second location (B), specifying the distance between the two locations (see Figure 1b).

The *what-direction* game is based on naming three locations: As in the *how-far* game, both agents are located at the first location (A) and they talk about the second location (B). The agents are both facing the third location (C), and the direction between the two distant locations is specified (see Figure 1c).

The *where-is-there* game, adapted from previous spatial language games (Bodík & Takáč, 2003; Steels, 1995), extends the *how-far* and *what-direction* games and is based on naming three locations, as specified in the *what-direction* game (see Figure 1d). The agents describe the relationship between the locations with spatial words of distance and direction. The *where-is-there* game is interesting because it allows the grounding of toponyms relative to existing toponyms, and therefore allows agents to refer to places that they have never visited or can never visit.

Cognitive Map

To build a representation of the world, the simulated robots used RatSLAM, a method of Simultaneous Localization And Mapping (SLAM) that has been developed over the past decade to enable autonomous robots to explore and map their environments (Milford & Wyeth, 2007). RatSLAM is a computational model inspired by the rodent hippocampal complex. Through exploration of an environment, each robot constructs a unique representation of the world as a topological map of experiences, each with an estimate of global pose within an approximate *x-y* representation of the world. An active experience encodes the robot's best estimate of its position (for more information see Milford, Schulz, Prasser, Wyeth, & Wiles, 2007). The experience map provides a cognitive map representation of the world (O'Keefe & Nadel, 1978).

A simulation world was built to mirror the real world, with images from the real world used in constructing the views of the robot. The simulation world includes an open plan office in a university building. Exploration was performed by left and right wall following. The robots used a single forward facing camera. In real-world studies, language games between real robots were based on actual hearing distances (Schulz, Wyeth, & Wiles, submitted). The study in this paper was completed in the simulation world for computational tractability. The simulation world enables simulated robots to pass messages to other robots within a set distance of their current locations, allowing the hearing distance to be explicitly set. For the study reported here a hearing distance of 3m was used.

Toponymic Lexicon

The associations between experiences and words are stored in distributed lexicon tables, a method inspired by the distributed nature of inputs to neural networks combined with the lexicon table structure (Schulz, et al., 2008). Forming concepts with a distributed lexicon table differs from most other conceptualization methods in that it is directly linked to the language formation, allowing concepts and words to have boundaries that are not explicitly defined. In many language game studies, concepts are formed using discrimination trees (Bodík & Takáč, 2003; Smith, 2007; Steels, 1997), which allows the agents to form concepts with well defined boundaries. The discrete concepts, formed through a discrimination tree or similar categorization method, may then be associated with words through a lexicon table. With a distributed lexicon table, concept formation and association with words occurs concurrently by increasing associations between experiences and words. An association value is stored for each experience–word pair, which is a value of 0.0 or greater. Experiences are related to each other by their proximity, based on their global pose estimates. The association between an experience and a word is strengthened when they are used together.

The toponymic lexicon data structures include the toponym lexicon, the toponym lexicon table, and toponym associations. The toponym lexicon comprises the set of words used as toponyms where each word is a unique string of consonants and vowels. The toponym lexicon table comprises a set of toponym associations between experiences and words.

In both the *where-are-we* and *where-is-there* games, the toponym association value for the specified experience and the word used is incremented by 1.0. A word for a location is chosen by the speaker in both the *where-are-we* and *where-is-there* games. For a specified location the word with the highest confidence value is chosen. The confidence value, h_{ij}, at the experience, i, for the word, j, is the relative association of the word within a neighborhood of size D compared to the total association of the word, calculated as follows:

$$h_{ij} = \frac{\sum_{k=1}^{X} a_{kj}^{T} (D^{T} - dist_{ki}^{T}) / D^{T}}{\sum_{m=1}^{E} a_{mj}^{T}} \qquad (1)$$

where X is the number of experiences within D of the experience, i; a^{T}_{ij} is the association between an experience, i, and the word, j; $dist^{T}_{ki}$ is the distance between experiences, k and i within the experience map of the robot; and E is the total number of experiences in the robot's experience map. For the study presented here a neighborhood size, D, of 3m was used. In each interaction, words are invented with probability, p, as follows:

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

where h_{ij} is the confidence value of the experience-word combination; and T is a scaling parameter called the temperature, which effectively sets the invention rate for new words. Eq. 2 allows agents to use existing words when a word is associated with the current location with a high confidence, and to probabilistically invent words otherwise. Varying the temperature alters the rate of word invention, where a higher temperature increases the probability of inventing a new word. For the study presented here the temperature was decreased linearly from 0.3 to 0.1 over the course of each generation.

Relational Lexicon

In addition to locations, the simulated robots have words for directions and distances. The data structures include the distance and direction lexicons, elements, associations, and lexicon tables. The distance lexicon comprises the set of words used to refer to distances, and the distance lexicon table comprises a set of distance associations between distance elements and words. Each distance element is a distance measured in meters in global pose space.

Direction words used data structures similar to those for distance words. The direction lexicon comprises the set of words used to refer to directions (i.e. angular distances), and the direction lexicon table comprises a set of direction associations between direction elements and words. Each direction element is an angle measured in radians.

In each *how-far* game, the association values stored in the distance lexicon for the distance word used are updated. Experiences are grouped to the nearest distance element based on their distance from the current experience in global pose space. For the topic, *j*, of the interaction, a distance association value, a^D_{ij}, is calculated for each distance element, $i \in 1..K^D$, by summing the target toponym associations for each experience grouped to that distance element, and smoothing using a distance neighborhood, as follows:

$$a_{ij}^{D} = \sum_{m=1}^{Y} \frac{\left(\sum_{k=1}^{X} a_{kw}^{T}\right)\left(D^{D} - dist_{mi}^{D}\right)}{D^{D}} \quad (3)$$

where Y is the number of distance elements within a neighborhood of size D^D from the distance element, *i*; X is the number of experiences grouped to the distance element, *i*; a^T_{kw} is the toponym association between the experience, *k*, and the toponym, *w*; and $dist^D_{mi}$ is the distance between the two distance elements, *m* and *i*. For the studies reported here, 50 distance elements were used in the range 0 to 25m and a distance neighborhood of 1.5m was used.

In each *what-direction* game, the association values stored in the direction lexicon for the direction word used are updated. Experiences are grouped to the nearest direction element based on the direction from the agent's facing at the current experience. For the topic, *j*, of the interaction, a direction association value, a^Θ_{ij}, is calculated for each direction element, $i \in 1..K^\Theta$, by summing the target toponym associations for each experience grouped to that direction element, and smoothing using a direction neighborhood, as follows:

$$a_{ij}^{\Theta} = \sum_{m=1}^{Y} \frac{\left(\sum_{k=1}^{X} a_{kw}^{T}\right)\left(D^{\Theta} - dist_{mi}^{\Theta}\right)}{D^{\Theta}} \quad (4)$$

where Y is the number of direction elements within a neighborhood of size D^Θ from the direction element, *i*; X is the number of experiences grouped to the direction element, *i*; a^T_{kj} is the toponym association between the experience, *k*, and the toponym, *w*; and $dist^\Theta_{mi}$ is the angular distance between the two direction elements, *m* and *i*. For the studies reported here, 50 direction elements were used in the range 0 to 2π, and a direction neighborhood of 3π/25 (21.6°) was used.

For distances and directions, the word with the closest match to the current distance or direction concept is used. The probability of inventing spatial words is calculated as for the toponyms using the match, $match_{ij}$, between the normalized vectors of the calculated, *i*, and stored, *j*, spatial associations, in place of the confidence value, calculated as follows:

$$match_{ij} = \sum_{k=1}^{K} \min\left(\frac{a_{ki}}{\sum_{m=1}^{K} a_{mi}}, \frac{a_{kj}}{\sum_{m=1}^{K} a_{mj}} \right) \quad (5)$$

where K is the number of spatial elements; a_{ki} is the association for the spatial element, *k*, and the topic, *i*, calculated using Equation 3 or 4; and a_{kj} is the association stored in the lexicon table for the spatial element, *k*, and spatial word, *j*.

Evolving Spatial Languages

In the study described in this paper, agent populations evolved languages over generations of agents. Generations consisted of a set number of interactions. In the initial population two agents played negotiation games. In subsequent generations, the older agent was replaced by a new agent. The new agent was the hearer (student) in all language games. When the new agent replaced the older agent in the following generation, all language games were played as the speaker (teacher). Note that the agents do not have fitness awarded and do not compete to be part of the next generation. There are always two agents per generation, with the older agent coming from the previous generation and the younger agent forming the next generation. In this view of language change, evolution refers to the change in the language rather than to the agents. Note that this use of evolution is consistent with its original Darwinian meaning as "descent with modification". Language change under this definition does not require direct competition of elements, rather it requires generations through which it is propagated, with features of the language affected by the generational transmission process.

The order in which concepts are formed by the agent can be constrained by the games played by the agent and the concepts chosen to be used in each game. In this study, the agents play *where-are-we* games initially to allow the separate formation of a set of toponyms then play *how-far* and *what-direction* games to form a set of relational terms and finally play *where-is-there* games. Agents play *where-are-we* games in all of the interactions of the generation, playing only *where-are-we* games for the first half of the interactions. In the third quarter of the interactions, agents may also play *how-far* and *what-*

direction games with equal probability, with the constraint that the agent must have at least two toponyms in order to play a *how-far* game and at least three toponyms to play a *what-direction* game. In the final quarter of interactions, the agents may also play *where-is-there* games, with the constraint that the agent must have at least one distance and one direction word.

The Language Bottleneck

The language transmission bottleneck refers to limited transmission of a language between generations. During its lifetime, a student may not be exposed to the entire lexicon of its teacher, or even when exposed to words, will learn its own grounded meaning and therefore will not perfectly learn the teacher's language. In this simulation study, the language bottleneck is due to limits on both the number of interactions per generation, and also the number of locations in the world where the agents interact. The student must therefore generalize from its experience of the teacher's language. How well the student can generalize depends on the number of interactions and the distribution of locations at which the interactions take place. The number of interactions per generation determines the proportion of the teacher's language that the student experiences during its lifetime. An initial investigation was performed with nine conditions based on 4, 8, 16, 32, 64, 128, 250, 500, and 1000 interactions per generation. The study comprised three runs of each condition with 20 generations per run.

The size of the language increased as the number of interactions per generation increased (see Table 1). The size of each lexicon differed, with larger toponym lexicons and smaller distance lexicons for more than 16 interactions per generation. For 4, 8, and 16 interactions per generation the direction lexicon was the smallest of the three lexicons. For each of the types of words (toponyms, distances, and directions), there was a crossover between more words invented per generation and more words retained per generation (see Figure 2). The crossover point indicates the number of interactions per generation where the language transmission bottleneck is sufficiently wide that more words are retained than invented. If a student learns a comprehensive language from its teacher, then proportionately fewer words will need to be invented in the next generation.

Language Change across Generations

For the conditions in which more words were preserved than invented, the language change can be investigated further. The three conditions considered further were a) 250, b) 500, and c) 1000 interactions per generation. The study comprised three runs of each condition to 20,000 interactions, consisting of a) 80, b) 40, and c) 20 generations.

In all three conditions, the simulated robots formed a shared set of toponyms, distances, and directions. The number of words in the lexicon of each agent for each type of word increased rapidly over the first few generations, and agents in all conditions continued to invent words for toponyms, distances, and directions throughout their lifetimes. The invention of words occurred at different rates in each

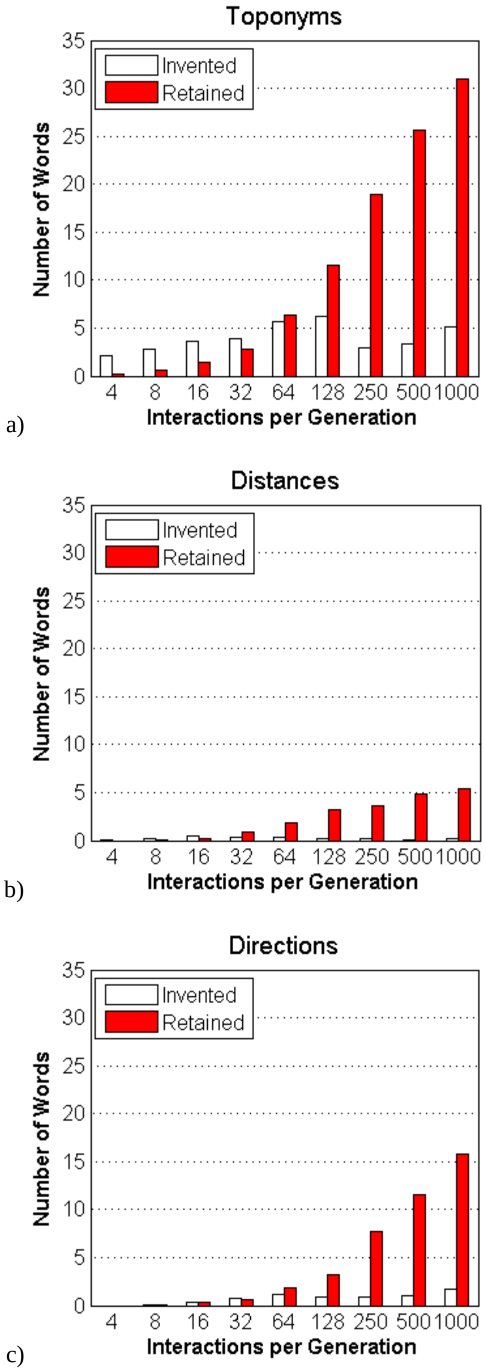

Figure 2. Words invented and retained per generation for each condition (4, 8, 16, 32, 64, 128, 250, 500, and 1000 interactions per generation) for a) Toponyms, b) Distances, and c) Directions. For each condition, the number of words invented and retained in each generation was averaged over the final ten generations of the three runs. Note the crossover between more words invented and more words retained occurs between 32-64 interactions per generation for toponyms (a) and directions (c), and 16-32 interactions per generation for distances (b).

condition and concept type (see Figure 3), with word loss closely matching word invention after the initial spurt of invention. The persistence of words in the lexicon through generations can be measured by considering when the words used in the final generation were initially invented. If a large proportion of the words were invented in earlier generations, then the words are persistent and the lexicon is stable. The persistence of words varied over the conditions and the concept types (see Figure 4).

Table 1. Average toponym, distance, and direction lexicon size over generations 11 to 20

Interactions per generation	Lexicon Size (mean (standard deviation))		
	Toponym	**Distance**	**Direction**
4	4.6 (1.1)	0.3 (0.5)	0.0 (0.0)
8	6.3 (1.5)	0.6 (0.8)	0.1 (0.4)
16	8.6 (2.2)	1.3 (0.8)	1.1 (0.7)
32	10.6 (1.7)	1.8 (0.6)	2.1 (1.1)
64	17.6 (2.8)	2.5 (0.8)	4.2 (1.2)
128	23.5 (4.4)	3.7 (0.5)	5.4 (1.5)
250	24.1 (3.1)	4.0 (0.7)	9.4 (1.3)
500	31.9 (3.2)	5.0 (0.5)	13.6 (1.7)
1000	40.4 (5.9)	5.8 (0.7)	19.3 (2.5)

Discussion

Learning with culture is different to inventing language from scratch. Agents begin their lives by learning words from older agents, and can later choose to use these words or invent new words. As agents start afresh in every generation, words that are no longer used do not remain in the lexicon. A change in language over time where one word or structure replaces another does not mean that the original is directly replaced by its replacement. Rather there may be an intermediate state in which either the old or the new word or structure may be chosen (Brinton & Traugott, 2005). In the studies presented here, an agent can learn a word for a location, but probabilistically also can invent a new word for the same location, while retaining representations for the old word.

The results show that a major effect of the length of the learning period was on the size of the resulting lexicon for the toponyms and the simple prepositions of distances and directions. The number of words used increased with the number of interactions per generation, as each agent had more interactions in which to learn the existing lexicon and invent new words. With shorter generations, the agents do not play a sufficient number of language games for a stable shared language to emerge.

The size of each lexicon is due to several factors, including the space of possible concepts, the neighborhood size used when choosing the appropriate word, the temperature used to set the probability of word invention and the opportunities to use words from that lexicon. The space of possible concepts is the size of the world for location and distance concepts and all directions for direction concepts. The neighborhood size for each word type is currently set to 3m for location concepts, 1.5m for distance concepts and 3π/25 (21.6°) for direction concepts. The opportunities to use the words are in the number of games of each type played.

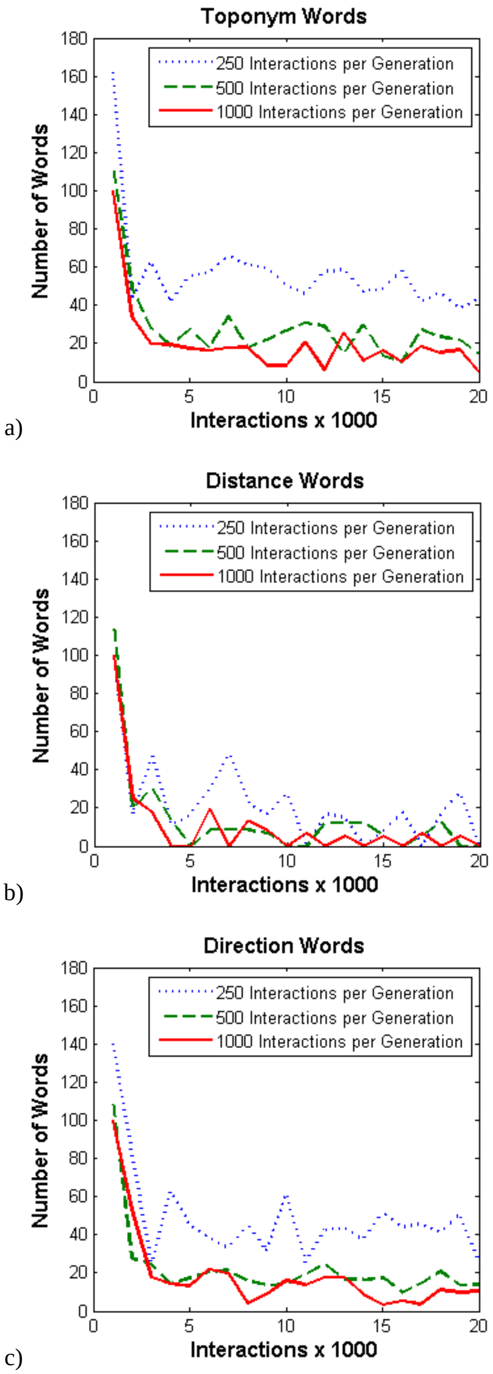

Figure 3. Words invented per 1000 interactions for the three conditions for a) toponyms, b) distances, and c) directions, averaged over all runs for each condition. In all conditions the word invention rate began high as the agent's lexicons developed over the first few generations. Distance words were more stable than direction words and toponyms, with fewer words invented and lost in each generation. The word invention rate for each type of word stabilized at a higher rate with a smaller number of interactions per generation.

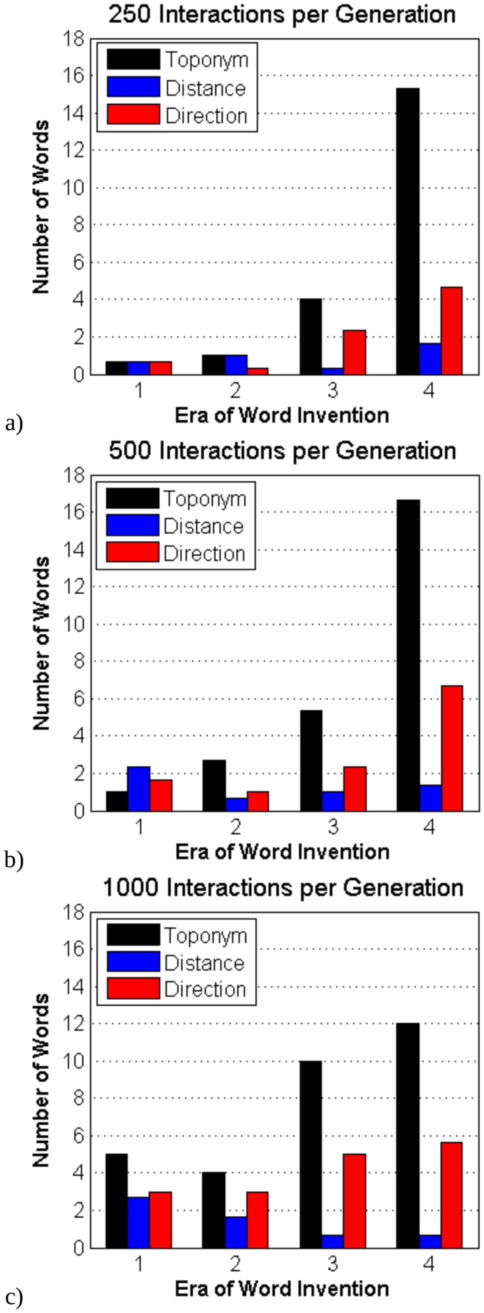

Figure 4. Word age in the final generation. The words used in the final generation are clustered into four eras based on the interaction in which each word was first used: 1. the early era (interactions 1 to 5,000), 2. the early-middle era (interactions 5,001 to 10,000), 3. the middle-late era (interactions 10,001 to 15,000), and 4. the late era (interactions 15,001 to 20,000). (a) For 250 interactions per generation few words were retained from earlier generations. (b) For 500 interactions per generation a higher proportion of distance words were retained from earlier generations. (c) For 1000 interactions per generation as well as retaining a higher proportion of distance words from earlier generations, the direction words in the final generation were invented more evenly across the generations, and a higher proportion of toponym words were invented in later generations.

With small numbers of interactions per generation, the small size of each lexicon is due to insufficient opportunities to play games that involve all possible locations. With larger numbers of interactions per generation, there is a trend towards a large toponym lexicon and a small distance lexicon. Smaller lexicons form when there is no noise in transmission and therefore no concepts that cover the same region in concept space. Larger lexicons form when the full concept space is covered. The main reason for the small size of the distance lexicon is likely to be that the size of the world has constrained the possible distances referred to by the agents. Increasing the size of the directly experienced world would result in the formation of a greater number of location and distance concepts. Direction concepts are restricted to one full rotation.

As shown by Smith (2007) and Bodík & Takáč (2003) a stable shared language can emerge in each longer generation, but the meaning of words may shift over generations, with new words entering the lexicon and old words forgotten. Bodík & Takáč (2003) found that more specific terms change faster than more general terms. If words enter and leave a language stochastically, the effect of the bottleneck would be the same for different classes of words. An alternative hypothesis is that unambiguous or frequently used words would pass through the bottleneck more easily than ambiguous or infrequent words. In the studies, we found differential rates of transmission for different classes of concepts, and saw the influence of the language transmission bottleneck on languages formed in conditions with both small and large numbers of interactions per generation.

The distance words were found to be more stable throughout the generations than the direction words. The stability of the words may be due in part to the smaller size of the distance lexicon. However, we conjecture that an equally important reason for more stable distance words is that compared to direction words, the creation of distance words is less noisy with only two toponyms used rather than three, and therefore their use is more reliable.

For the conditions explored in this study, in which word retention is higher than word invention, the bottleneck of language transmission is still evident in the trends for word age across the conditions and types of words. Proportionately more words were invented in later generations for all condition and concept types except for distance words in the conditions of 500 and 1000 interactions per generation. In these conditions, the early distance words pass through the bottleneck unchanged. In all other conditions and word types, the language transmission bottleneck reduces the retention of words through generations of agents.

As discussed in the introduction, a variety of factors have been identified as contributing to language change (for example, see Aitchison, 1991; Kirby, et al., 2007; Niyogi, 2006). Some factors contributing to language change have been demonstrated in the studies presented here. The size of the lexicon was affected by the social interactions and the period of individual language learning, and the rate of change for different concept types was affected by the concept formation for each word type. We have shown that learning periods and concept formation affect the rate at which words are retained, invented, and lost from the lexicon of the agent population. The key contribution of this research is a demonstration of the impact of language acquisition (in the form of individual language learning, concept formation, and social interactions) on language change, in particular showing that the bottleneck of language transmission can still affect word retention between generations even when a stable shared language forms within each generation.

Acknowledgements

We thank members of the RatSLAM team and colleagues in the Thinking Systems project. This research was supported by Australian Research Council grants to GW and JW.

References

Aitchison, J. (1991). Language change: progress or decay? (2nd ed.). Cambridge: Cambridge University Press.

Bodík, P., & Takáč, 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.), Frontiers in Artificial Intelligence and Applications: Eighth Scandinavian Conference on Artificial Intelligence SCAI'03 (pp. 37-46). Amsterdam: IOS Press Inc.

Brighton, H., Smith, K., & Kirby, S. (2005). Language as an evolutionary system. Physics of Life Reviews, 2(3), 177-226.

Brinton, L. J., & Traugott, E. C. (2005). Lexicalization and language change. Cambridge, UK: Cambridge University Press.

Chater, N., & Christiansen, M. H. (2009). Language Acquisition Meets Language Evolution. Cognitive Science.

Kirby, S. (2002). Learning, bottlenecks and the evolution of recursive syntax. In E. J. Briscoe (Ed.), Linguistic Evolution through Language Acquisition: Formal and Computational Models: Cambridge University Press.

Kirby, S., Dowman, M., & Griffiths, T. L. (2007). Innateness and culture in the evolution of language. Proceedings of the National Academy of Sciences, 104(12), 5241-5245.

Kirby, S., & Hurford, J. R. (2002). The emergence of linguistic structure: an overview of the iterated learning model. In A. Cangelosi & D. Parisi (Eds.), Simulating the Evolution of Language (pp. 121-148). London: Springer Verlag.

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.

Milford, M., & Wyeth, G. (2007). Spatial mapping and map exploitation: A bio-inspired engineering perspective Spatial Information Theory (pp. 203-221). Berlin: Springer.

Niyogi, P. (2006). The computational nature of language learning and evolution. Cambridge, Massachusetts: MIT Press.

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

Schulz, R., Prasser, D., Stockwell, P., Wyeth, G., & Wiles, J. (2008). The formation, generative power, and evolution of toponyms: Grounding a spatial vocabulary in a cognitive map. In A. D. M. Smith, K. Smith & R. Ferrer i Cancho (Eds.), The Evolution of Language: Proceedings of the 7th International Conference (EVOLANG7) (pp. 267-274). Singapore: World Scientific Publishing.

Schulz, R., Wyeth, G., & Wiles, J. (submitted). Lingodroids: Grounding spatial language in a cognitive map.

Smith, A. D. M. (2007). Language change and the inference of meaning. In C. Lyon, C. Nehaniv & A. Cangelosi (Eds.), Emergence of Communication and Language (pp. 323-337). London: Springer.

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

Steels, L. (1997). The origins of syntax in visually grounded robotic agents. In M. Pollack (Ed.), Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence (IJCAI-97) (Vol. 2, pp. 1632-1641). San Francisco, California: Morgan Kauffman Publishers.

Tonkes, B., & Wiles, J. (2002). Methodological issues in simulating the emergence of language. In A. Wray (Ed.), The Transition to Language. Oxford: Oxford University Press.

Vogt, P. (2004). Minimum cost and the emergence of the Zipf-Mandelbrot law. In J. Pollack, M. A. Bedau, P. Husbands, T. Ikegami & R. A. Watson (Eds.), Artificial life IX: Proceedings of the Ninth International Conference on the Simulation and Synthesis of Living Systems: MIT Press.