

Stable communication through dynamic language

Andrew D. M. Smith*

*Language Evolution and Computation Research Unit,
School of Philosophy, Psychology and Language Sciences
University of Edinburgh
Adam Ferguson Building
40 George Square
Edinburgh, EH8 9LL
United Kingdom
andrew@ling.ed.ac.uk

Abstract

I use agent-based computational models of inferential language transmission to investigate the relationship between language change and the indeterminacy of meaning. I describe a model of communication and learning based on the inference of meaning through disambiguation across multiple contexts, which is then embedded within an iterated learning model. The dynamic flexibility and uncertainty inherent in the model leads directly to variation between agents, in both their conceptual and lexical structures. Over generations of repeated meaning inference, this variation leads to significant language change. Despite such change, however, the language maintains its utility as a communicative tool within each individual generation.

1 Introduction

All living human languages are constantly changing. Tiny, often barely perceptible, changes in the contexts in which particular words are used, or the way in which they are pronounced, accumulate over generations of use to such an extent that the language itself becomes unrecognisable in only a few generations. The driving force behind historical linguistic change is widely recognised to be linguistic variation (Trask, 1996). In this paper, I explore the relationship between the incessancy of language change and the indeterminacy of meaning, using an agent-based computational model of iterated inferential communication. Inferential communication focuses on the fact that information is not transferred directly between communicants, but rather indirectly: the hearer infers the meaning of a signal from pragmatic insights and the context in which the signal is heard. The uncertainty inherent in this process means that individuals do not necessarily infer the same meanings, leading to differences in their internal linguistic representations. Over generations of inferential communication, small variations may result in significant levels of language change.

Inferential models of language have already been used successfully to model learning conceptual structures and language in tandem (Smith, 2003b) and

the effects of psychologically plausible constraints on lexical acquisition (Smith, 2005), but have not yet been used in the detailed study of process of language change. In the experiments presented here, I build on the basic inferential model, embedding it within a successful model of repeated cultural transmission with generational turnover (Smith et al., 2003), in order to explore the nature and extent of language variation and change across many generations of language users.

The remainder of this paper is divided into five parts. In section 2, I explore the twin theoretical foundations on which the model used to perform the experiments is based, namely cultural transmission and the inference of meaning. In section 3, I describe the model in detail, including how agents create meanings, communicate with each other, and infer meaning from multiple contexts. In section 4, I discuss the different kinds of variation which are present in the model, and how these can be measured. In section 5, I present the results of the experiments themselves, and demonstrate that conceptual and lexical variations result in remarkably rapid and significant change to the language itself. At the same time, the language's utility as a successful shared communication system is reconfirmed within each generation. Finally, in section 6, I provide a summary of the paper's main conclusions.

2 Foundations

2.1 Cultural Transmission

Although we are all genetically endowed with the cognitive capacity to learn and use language, the particular languages we actually learn are not stored in our genes, but are instead those which we hear spoken by people in the communities in which we live. Languages are therefore passed on culturally. Recent research into language evolution has focused on this cultural nature of transmission, building models which represent the external and internal manifestations of language as distinct phases in the language's life cycle: individuals produce their *external* linguistic behaviour based on their internal linguistic representations, and in turn induce their own *internal* linguistic representations, or grammars, in response to the linguistic behaviour, or primary linguistic data, which they encounter. Such models of linguistic evolution are known as expression/induction (E/I) models (Hurford, 2002) or iterated learning models (Smith et al., 2003). The cultural nature of these models is captured in the fact that the linguistic input used by one individual to construct its grammar is itself the linguistic output of other individuals. Differences which occur between the internal grammars of individual members of the population occur as a result of the dynamic cultural evolution of the language itself.

Iterated learning models have been used successfully to demonstrate the cultural emergence of a number of structural characteristics of language, notably compositionality (Brighton, 2002) and recursion (Kirby, 2002). These properties arise through repeated cultural transmission when agents must learn a language made up of signal-meaning mappings from a restricted set of data, through a transmission *bottleneck*. Under such conditions, holistic, idiosyncratic rules of grammar can only be successfully transmitted if the specific signal-meaning pair is encountered. Compositional rules, on the other hand, are preferentially produced due to their generalisability, and thus are much more likely to pass through the bottleneck into the next generation (Smith et al., 2003).

2.2 Meaning Inference

It is important to note, however, that many such models of cultural evolution are characterised by the explicit coupling of pairs of signals and predefined meanings. This coupling necessarily leads to the development of syntactic structure which is identical to the predefined semantic structure, and which undermines, to a significant extent, the claims for emer-




Figure 1: A model of communication which avoids the signal redundancy paradox. The model has three levels of representation: an external environment (A); an internal semantic representation (B); and a public set of signals (C). The mappings between A and B and between B and C, represented by the arrows, fall into the internal, private domain, whose boundary is shown by the dotted line.

gence. In these models, a linguistic utterance consists of the explicit conjunction of a signal and a meaning, and communication involves the direct transfer of this utterance between agents. In communication, then, both the signal and the meaning are simultaneously transferred. As I have shown previously, however, if the meanings are directly transferred, then there is no role for the signals to play, leading to the *paradox of signal redundancy* (Smith, 2003b, 2005): what is the motivation for language users to spend time and energy in learning a symbolic system of signals which provides them no information that they do not already have from the directly transferred meanings?

The inferential model presented here is motivated to a large extent by avoiding the signal redundancy paradox. This is easily done by recognising that meanings are not directly transferable. Instead, a meaning is encoded into a signal by the speaker, and decoded back by the hearer. Of course, decoupling meanings and signals means that there is now no easy way for agents to associate them with each other, and so we must assume that meanings are *inferred* from some external source. The mere existence of an external world is not sufficient to avoid the signal redundancy paradox, however; we must also insist on a strong demarcation in the model between the external world and the agents' internal representations, as shown in figure 1. The external, or public, domain contains objects and situations which can be potentially accessed and manipulated by all agents, while

the internal, or private, domains are accessible only by a particular agent, and contain representations and mappings created and developed by the agent itself.

Signals and their referents are linked only indirectly, mediated via separate associative mappings between themselves and each agent's internal meaning representations. The associative mappings themselves, however, are created individually by each agent through analysis of the co-occurrence of signals and referents over multiple situations, as described in section 3.2 below.

3 The Inferential Model

The E/I models of cultural transmission described in this article, therefore, contain neither a predefined, structured meaning system, nor an explicit link between signals and meanings. Instead, I describe experiments with simulated agents who initially have neither conceptual nor lexical structures, but have the ability to *create* conceptual representations and to *infer* meaning from their experiences. The model contains an external world with a number of objects, which can be objectively described in terms of the values of their abstract features, real numbers generated within the range [0:1]. Agents are provided with dedicated sensory channels, which they can use to sense whether a particular feature value falls within two bounds, and use these to create meanings which allow them to distinguish objects from each other. Agents can also create words to express these meanings and to communicate about the objects. This model is based on that described initially by Steels (1996), in which two agents (a speaker and a hearer) play a series of language games, but is extended in a number of ways. In the following sections, I describe how agents create meanings in response to their interactions with the external world, how they create and use signals to communicate to each other about situations in the world, and how they infer the meanings of signals they receive. Finally, I explain how the inferential model is placed within an iterated learning paradigm, to allow experiments exploring the nature and extent of language change across generations.

3.1 Meaning Creation

Meaning creation occurs as agents explore their environment and try to discriminate objects from each other. In such an exploratory episode, an agent investigates a random subset of objects, called the *context*, with the aim of distinguishing one particular, randomly-chosen *target* object within the context

from all the other objects therein. The agent searches its sensory channels for a *distinctive category*, an internal semantic representation which accurately describes the target, but does not accurately describe the other objects in the context. If no such category exists, and so the episode fails, the agent expands its semantic capacity, by splitting the sensitivity range of an existing category into two halves, thereby creating two new categories. Repeated meaning creation in this way results in the development of hierarchical, tree-like conceptual structures where the nodes on the tree represent semantic categories. Nodes nearer the tree root represent more general meanings, with wider sensitivity ranges which cover a greater proportion of the semantic space, while those nearer the leaves represent more specific meanings. Importantly, the simulations contain no pre-specification of which categories should be created, and meaning creation is carried out by each agent individually according to its own experiences. This means that individual agents create different, but typically equally valid, conceptual representations of their world.

3.2 Inferential Communication

Communication follows from a successful discrimination episode. Having found a distinctive category, the speaker chooses a suitable signal from its lexicon to represent it; if none is appropriate, then the speaker creates a new signal as a random string of letters. The signal is then transmitted to the hearer, who also observes the original context from which the speaker derived its distinctive category. Importantly, however, neither the distinctive category nor the target object to which it refers are ever identified to the hearer.

Hurford (1989) developed dynamic communication matrices of transmission and reception behaviour to model the evolution of communication strategies, and showed that bidirectional, Saussurean mappings between signals and meanings are essential in the development of viable communication systems. Oliphant and Batali (1997) extended this model to show that the best way to ensure continuing increases in communicative accuracy is for speakers to always choose signals based on how they are interpreted by the rest of the population. Their algorithm, however, requires agents to be able to have direct access to the internal representations of other agents. In order to avoid this mind-reading, I have used a modified version of the algorithm, *introspective obverter* (Smith, 2003b), in which the speaker chooses the signal which it would be most likely to interpret correctly, given the current context. Because the speaker

cannot access the interpretative behaviour of the other agents, signal choice is based on the speaker's own interpretative behaviour.

Once provided with a signal, but without any information about the meaning or the object to which it refers, the hearer must infer the meaning from the information in the context, and from its previous experience of the signal in other contexts. Inference takes place through *cross-situational statistical learning* (Smith and Vogt, 2004). In every situation in which a word is encountered, the hearer creates a list of *semantic hypotheses*, or every possible meaning which could serve as a distinctive category for *any* single object in the current context. Each of these meanings is then associated with the signal in the hearer's internal lexicon. The lexicon contains a count of the co-occurrence of each signal-meaning pair $\langle s, m \rangle$, which is used to calculate the conditional probability that, given s , m is associated with s . The hearer simply chooses the meaning with the highest conditional probability for the signal it receives and assumes that this was the intended meaning.

If the hearer's chosen meaning identifies the same object as the speaker's initial target object, then the communicative episode is deemed successful. Communicative success is therefore based on *referent identity*: there is no requirement for the agents to use (or even to have) the same internal meaning, but they must identify the same external referent. Furthermore, neither agent receives any feedback about the communicative success of the episode, so the only information available for use in the inferential process is the co-occurrence of signals and referents across multiple contexts. This method of cross-situational inferential learning is similar to the method proposed by Siskind (1996), but differs from it most fundamentally in that the set of possible meanings over which inferences are made in the model presented here is neither fixed nor predefined, but is instead dynamic, and in principle infinite.

Previous experiments using cross-situational statistical learning show that the method is powerful enough for agents to learn large lexicons, and that agents with different conceptual structures can communicate successfully. The time taken to learn a whole lexicon is primarily dependent on the size of the context in which each item is presented (Smith, 2003a; Smith and Vogt, 2004), while communicative success is closely related to the level of inter-agent meaning similarity (Smith, 2003b). However, if agents are endowed with psychologically motivated interpretational biases to aid inference, such as mutual exclusivity (Markman, 1989), then even agents

with very dissimilar conceptual structures can communicate successfully (Smith, 2005).

3.3 Iterated Inference

In order to explore how languages change over generations, the inferential model is then extended vertically into a traditional iterated learning model with generational turnover (Smith et al., 2003). It is helpful in this regard to consider the speaker as an adult, and the hearer as a child. Each generation consists of a number of exploratory episodes, in which both agents explore the world individually and create meanings to represent what they find, followed by a number of communicative episodes, in which the adult communicates to the child. At the end of a generation, the adult is removed from the population, the child becomes an adult, and a new child is introduced. The language which was inferred in the previous generation by the child becomes the source of its own linguistic output in the subsequent generation, when it is an adult. This process of generational turnover is then iterated a specified number of times.

4 Variation

It is well recognised that language change is driven by various kinds of variation in language communities (Trask, 1996). In the inferential model I have sketched above, there are two main sources of variation, which I will call *conceptual* and *lexical*. In the following sections, I will describe the source and effects of both types of variation, examples of which can be seen in figure 2. Taken from a representative simulation, this shows an extract from the conceptual and lexical structure of an adult and a child from the same generation. Each agent actually has five sensory channels on which conceptual structures are built, but only one of these channels is shown in figure 2.

4.1 Conceptual Variation

The independent creation of conceptual structure based on individual experience leads inevitably to variation in agents' conceptual representations, both because an agent's response to a certain situation is not deterministic, and because the experiences themselves differ between agents. The relative similarity of two agents' conceptual representations can be quantified by measuring the tree structures built on each sensory channel, then averaging across each sensory channel (Smith, 2003a). If $k(t, u)$ is the number of nodes which two trees t and u have in common,




Figure 2: Extract from the internal structures of two agents, showing variation in both conceptual and lexical structures. The conceptual structures are shown by hierarchical tree structures, each node of which represents a different meaning. Conceptual variation, where meanings have no corresponding equivalent in the other agent’s conceptual structure, is marked with dotted lines and colour. Lexical structures are represented by the words attached to the nodes, which signify the agent’s preferred word for the meaning; empty nodes have no preferred word. Lexical variations, where the agents disagree on the meaning of a word, are circled.

and $n(t)$ is the total number of nodes on tree t , then the similarity $\tau(t, u)$ between trees t and u is:

$$\tau(t, u) = \frac{2k(t, u)}{n(t) + n(u)}.$$

By averaging this measure across all sensory channels, we can produce an agent-level measure of overall conceptual similarity. If a_{ij} identifies the tree on channel j for agent i , and each agent has c sensory channels on which they develop conceptual structure, then the meaning similarity $\sigma(a_1, a_2)$ between agents a_1 and a_2 is:

$$\sigma(a_1, a_2) = \frac{1}{c} \left(\sum_{j=0}^{c-1} \tau(a_{1j}, a_{2j}) \right).$$

In figure 2 above, we consider only the nodes on the trees themselves, without reference to the words attached to them. Nodes which have no equivalent in the other agent’s conceptual structure are marked with dotted lines and colour. This shows clearly that, although similar, the agents have developed different tree structures: the child has created additional conceptual structure in three different places.

4.2 Lexical Variation

The inherent uncertainty in the process of meaning inference through cross-situational learning also produces variations in the lexical associations made by

the agents. Not only are the inferred meanings dependent on the particular conceptual structure the hearer has created, but the associations themselves depend on the particular contexts in which the words are heard. Lexical variation can be measured by considering whether two agents have the same preferred word for any given meaning. An agent’s set of preferred words is calculated by sorting a copy of its entire lexicon in descending order of conditional probability (see section 3.2), then manipulating as follows:

1. find the topmost lexical entry, which is made up of signal s and meaning m .
 - (a) store s as the preferred word for m ;
 - (b) delete all lexical entries containing s .
2. repeat step 1, until the lexicon is empty.

In figure 2, preferred words are represented by the words attached to the appropriate nodes on the tree structure; empty nodes have no preferred word. If adult and child both have the same preferred word for a meaning, then the child has successfully learnt the word, and the lexical item has *persisted* through the generation. Lexical items which do not persist have undergone different kinds of semantic change; these are shown as circled words in figure 2. For example, the words wm and hhd have not been learnt successfully, despite the relevant nodes in the adult’s conceptual structure also existing in the child’s structure. In both of these cases, the words are associated

with nodes nearer the root of the child’s tree than the adult’s; because nodes nearer the root of a tree cover a larger area of semantic space, I consider this kind of change as a generalisation. Other kinds of semantic change, such as specialisation and analogy are not discussed further here.

Lexical persistence across the whole of an agent’s lexicon is very useful as a broad measure of linguistic change, and can be measured both within and between generations: *intra-generational* lexical persistence is the proportion of the adult’s lexicon learnt by the child, while *inter-generational* lexical persistence is the proportion of the original language developed by the adult in the first generation of the simulation which is still intact in the language of the child at the end of the n th generation.

5 Experimental Results

The aim of these experiments was twofold. Firstly, I wanted to verify whether results obtained in previous experiments with an inferential model in a single generation, briefly summarised in section 3.2, would remain valid in a multi-generational model. More importantly, I wanted to measure how languages themselves change over a number of generations, to explore whether languages undergoing rapid language change over successive populations of language users could still be communicatively viable.

5.1 Communicative Success and Meaning Similarity

I have previously shown in mono-generational inferential models that levels of communicative success are closely correlated with levels of meaning similarity between agents (Smith, 2003b). Figure 3 shows results from a typical simulation run over ten generations, each of which is made up of 20,000 episodes. Analyses of communicative success and meaning similarity were calculated every 1000 episodes: communicative success measures the proportion of successful communications over the previous 1000 episodes, while meaning similarity is measured as described in section 4.1.

We can clearly see that levels of meaning similarity and communicative success are again very closely correlated, as expected. In each generation, the communicative success rate rises rapidly at first, as the child successfully learns the meanings of many words through cross-situational inference. The rate then climbs more slowly, as the child tries to deduce the meanings of the remaining words. These represent




Figure 3: Communicative success and meaning similarity in an iterated inference model. Each generation consists of 20000 episodes.

meanings which are seldom used by the adult as distinctive categories, and so consequently occur relatively infrequently in communicative episodes, which makes the process of disambiguation through exposure in different contexts much slower.

Levels of communicative success and meaning similarity at the end of each generation were also measured, to see if any inter-generational trends were present, but we can clearly see in figure 3 that the levels of communicative success and meaning similarity achieved at the end of each generation were very similar, and in fact no significant inter-generational changes are discernible. This latter results contradicts recent work by Vogt (2003), however, who claims a small increase in inter-generational communicative success in simulations run through his Talking Heads simulator, using a similar model of inferential learning, which he calls selfish games.

5.2 Lexical Persistence

Secondly, I explored changes in the languages themselves over generations of different lengths, measuring lexical persistence to determine the nature and extent of change. Figure 4 shows both inter-generational and intra-generational lexical persistence over simulations of ten generations.

A comparison of the two graphs in figure 4 shows us how the length of a generation (the number of episodes which it contains) affects both measures of lexical persistence. Generations containing 5000 episodes (left) result in intra-generational lexical persistence rates at the end of each generation of between 60 and 70%, but if the generation length is increased to 20,000 episodes (right), then the lexical persistence




Figure 4: Inter-generational and intra-generation lexical persistence. Each generation consists of 5000 episodes (left) and 20,000 episodes (right).

rates are closer to 80%. Unsurprisingly, given more exposure to the language, the child is able to learn a higher proportion of it successfully. Note, however, that variation in the conceptual structures of the agents provides an effective ceiling for the level of intra-generational lexical persistence, as it is impossible for the child to learn the meaning of a word if the corresponding conceptual structure does not exist in its repertoire.

Figure 4 also shows that there are no significant differences between the levels of lexical persistence obtained within different generations. It is clear, on the other hand, that the rate of *inter-generational* lexical persistence shows a considerable cumulative decline after only a few generations. There are two separate pressures on the language which enforce its relentless erosion over successive generations of cultural transmission through inference, which can be regarded as twin bottlenecks on the language’s transmission.

1. Conceptual variation restricts the number of words which can *potentially* persist into the next generation: only words which refer to meanings which are shared are available to be learnt.
2. Lexical variation, or imperfections in inferential learning, further restricts the number of words which *actually* persist into the next generation.

The pressures from these two bottlenecks naturally result in a cumulative decline in inter-generational lexical persistence. These pressures are compounded in subsequent generations, so that even after only a few generations are passed, very little of the original language remains, and we find a language which

is changing very rapidly on an inter-generational timescale. Importantly, however, we can see from figure 3 that this rapid language change does *not* affect levels of communication within a single generation, which remain very high.

5.3 Generally Stable

If we investigate in more detail the languages which are used by the child at the end of each generation, we find that there is a distinct pattern to the language change which occurs. Words referring to more specific meanings tend to disappear first, and only more general words tend to survive across multiple generations. There are two obvious reasons for this, both artefacts of the design of the model. Firstly, the Steelsian method of hierarchical conceptual construction forces some order on the meanings which are created: there is no way, for instance, to create a meaning in the depths of a tree without first creating the relevant meanings further up the hierarchical structure. This restriction necessarily means that the more general meanings nearer the root of the tree, are more likely to be *shared* by the agents, and therefore less likely to be excluded from being learnt by the conceptual variation bottleneck. Secondly, agents use a model of communication which follows the maxim of quantity in Grice (1975)’s philosophical model of conversation, by choosing as distinctive categories meanings which provide sufficient information to identify the target object, but are not unnecessarily specific. This means in turn that more general meanings are more likely to be both used by the adult and also *inferred* by the child, and so are much more likely to pass through the second bottleneck on learning.

6 Conclusions

Although the cultural nature of language transmission is becoming more widely recognised, its inferential character is less widely acknowledged. Inferential communication not only provides an explanation for the existence of otherwise redundant signals, but also allows the construction of realistic models of dynamic language, in which uncertainty, variation and imperfect learning play crucial roles.

In this article, I have briefly presented a model of language as a culturally transmitted system of communication, based on the creation and inference of meaning from experience. Individual meaning creation, and the uncertainty inherent in meaning inference lead to different degrees of variation in both conceptual and lexical structure. Conceptual variation and imperfect learning create twin bottlenecks on transmission, which result in rapid language change across generations. Despite this rapid language change, however, within each single generation the language itself remains sufficiently stable to establish and maintain its utility as a successful communication system.

Acknowledgements

Andrew Smith was supported by ESRC Research Grant RES-000-22-0266, and by ESRC Postdoctoral Research Fellowship PTA-026-27-0094. He would like to thank Jim Hurford, Kenny Smith and Paul Vogt for their valuable input and comments.

References

- H. Brighton. Compositional syntax from cultural transmission. *Artificial Life*, 8(1):25–54, 2002.
- H. P. Grice. Logic and conversation. In P. Cole and J. L. Morgan, editors, *Syntax and Semantics*, volume 3, pages 41–58. Academic Press, New York, 1975.
- J. R. Hurford. Biological evolution of the Saussurean sign as a component of the language acquisition device. *Lingua*, 77:187–222, 1989.
- J. R. Hurford. Expression/induction models of language evolution: dimensions and issues. In E. Briscoe, editor, *Linguistic Evolution Through Language Acquisition: Formal and Computational Models*, pages 301–344. Cambridge University Press, Cambridge, 2002.
- S. Kirby. Learning, bottlenecks and the evolution of recursive syntax. In E. Briscoe, editor, *Linguistic Evolution through Language Acquisition: Formal and Computational Models*, pages 173–203. Cambridge University Press, Cambridge, 2002.
- E. M. Markman. *Categorization and naming in children: problems of induction*. Learning, Development and Conceptual Change. MIT Press, Cambridge, MA, 1989.
- M. Oliphant and J. Batali. Learning and the emergence of coordinated communication. *Center for Research on Language Newsletter*, 11(1), 1997.
- J. M. Siskind. A computational study of cross-situational techniques for learning word-to-meaning mappings. *Cognition*, 61:39–91, 1996.
- A. D. M. Smith. *Evolving Communication through the Inference of Meaning*. PhD thesis, Philosophy, Psychology and Language Sciences, University of Edinburgh, 2003a.
- A. D. M. Smith. Intelligent meaning creation in a clumpy world helps communication. *Artificial Life*, 9(2):175–190, 2003b.
- A. D. M. Smith. Mutual exclusivity: Communicative success despite conceptual divergence. In M. Tallerman, editor, *Language Origins: perspectives on evolution*, pages 372–388. Oxford University Press, Oxford, 2005.
- A. D. M. Smith and P. Vogt. Lexicon acquisition in an uncertain world. Paper given at the 5th International Conference on the Evolution of Language, Leipzig, 2004.
- K. Smith, H. Brighton, and S. Kirby. Complex systems in language evolution: the cultural emergence of compositional structure. *Advances in Complex Systems*, 6(4):537–558, 2003.
- L. Steels. Perceptually grounded meaning creation. In M. Tokoro, editor, *Proceedings of the International Conference on Multi-agent Systems*, Cambridge, MA, 1996. MIT Press.
- R. L. Trask. *Historical Linguistics*. Arnold, London, 1996.
- P. Vogt. Grounded lexicon formation without explicit reference transfer. In W. Banzhaf, T. Christaller, J. Ziegler, P. Dittrich, and J. T. Kim, editors, *Advances in Artificial Life: Proceedings of the 7th European Conference on Artificial Life*, pages 545–552. Springer-Verlag, Heidelberg, 2003.