

A Platform for Simulating Language Evolution

Carl Vogel and Justin Woods*

Computational Linguistics Group

Trinity College Dublin

Dublin 2, Ireland

vogel@tcd.ie

Abstract

A platform for conducting experiments in the simulation of natural language evolution is presented. The system is parameterized for independent specification of important features like: number of agents, communication attempt frequency, agent short term memory capacity, communicative urgency, etc. Representative experiments are demonstrated.

1 Background

A large and growing literature on natural language evolution argues specific claims (e.g. the bidirectional sign emerges as an optimal linguistic information structuring device (Hurford, 1989)), through construction of special purpose simulations, and demonstrating relevant emergent behavior within those systems. Any simulation of emergent behavior has to be scrutinized for the possibility that the behavior is actually determined by the way underlying assumptions are encoded in initial conditions of the system. If initial conditions, rather than parameters of influence, determine the emergent behavior, then it has not truly emerged. Equally, some of the encoded assumptions may not in fact determine the outcome, but may be nonetheless questionable and worthy setting as an input parameter rather than as a precondition. Many assumptions are innocuous, but others are fundamental to the sorts of results that can emerge in these tailored systems.¹ When these risks are attended to properly, a very interesting program of research can be pursued, as has been documented widely in the literature. An argument is presented here that a platform is necessary which is parameterized for as many of these assumptions as possible so that the interaction of values for the relevant parameters may be more comprehensively explored in a common setting. A system which meets some of these requirements is detailed in §2. The system has been under development for some time (Vogel & Woods, 2002). This article describes its present incarnation, and the results of experiments conducted within it in response to relevant questions that persist in the literature about language evolution.

*Now: European Patent Office, Gitschiner Strasse 103, D-10969 Berlin (jwoods@epo.org)

¹Steels (1997), for example simulates a world in which only 10 objects exist.

The simplifying assumption of semantic transparency is surprisingly prevalent in research practice in simulating language evolution. For example, Hurford (1989) simulated populations of interlocutors associating concepts with symbols on the basis of different sampling strategies. The sampling is of how other agents map symbols to concepts when they produce utterances and when they interpret utterances. In his simulations, the Saussureans, the ones who use one sort of sampling to dictate their own utterances and interpretations, are the winners. A bidirectional sign emerges from the assumption that it is possible to ascertain what other language users mean by something at the very start of language development. Kirby (1998) considers the extent to which language development (its complexities and systematicities) can be explained without direct genetic factors. He measures the size of the grammar in relation to the number of expressible meanings over time. He models utterances as pairs of forms and meanings. "Each individual learns only from utterances (form-meaning pairs) produced by its neighbors" (Kirby, 1998, p. 7). Thus, here again learners of the language have direct access to intended meaning.

Here as well, the assumption is troubling. Other researchers have also noted this. Oliphant (1999) describes a social/cultural model with a consideration of genetic factors. Smith (2002) discusses research on either side of the assumption that language is a genetic or social construct. Oliphant comments on the fact that it is problematic to assume that hearers are presented with transparent access to intended meanings, with an example isomorphic to Quine's "Gavagai". However, in the simulations he constructs he nonetheless assumes that learners have simultaneous access to signals and meanings, leaving it as "an open question how learners extract the signal/meaning pairs that they associate" (Oliphant, 1999, p.g. 16). The intention here is to examine directly this association without assuming it as a given. If early agents originally had perfect communication of meaning as a means to make arbitrary linguistic associations with them, there is nothing to motivate making the mapping. The only apparent advantage of language as a transparent representation system for meanings is in recording meanings in the context of a broken chain of communicators (if one has a non-broken chain of communicators, then meanings can be passed from one directly to another) and in that case the language as a mapping between arbitrary sounds and meanings might serve as a means of communicating otherwise lost information. However, this would suggest either that writing systems would be much older than they are or that in the early days of language, parrots were more prolific and interwoven into human life than researchers have until now concluded from fossil evidence. There is, of course, a lot of advantage in having language available as a distinctly imperfect representation system with peace-keeping potential inherent in ambiguity and other individual advantages that might accrue from being able to lie. However, an imperfect representation system would have little need to assume prior perfect communication of meaning to accompany the bootstrapping.

Another assumption implicit in the evolution literature which requires close examination is the speed with which especially social models of linguistic development should proceed towards convergence. For example, the work mentioned

by Oliphant (1999) includes experiments that reach a stable state after about 1000 rounds; similarly that of Steels (1997). However, particularly if language is a socially sustained construct, it would seem that for a language to come into existence at all, and to persist as a cultural artifact, it must prove quite useful quite quickly. Even assuming that it does make communication easier, thus enhancing adaptive value, one would want a situation in which without cheating by allowing interlocutors free access to meaning anyway roughly successful communication (or at least not obviously unsuccessful communication) emerges extremely quickly. It is extremely difficult to imagine the first users of language lasting more than a few dozen unsuccessful attempts at linguistic behavior, much less thousands.

Certain fundamental features of the simulation platform are still lacking, and arguably some will always be lacking: a simulation of natural language evolution will never count as the evolution of a natural language. However, the factors which impinge on the simulation may well be related to the real thing.

2 A Language Evolution Workbench

2.1 Design Issues

The experiments reported in §3 are based on a general architecture for simulating social models of language evolution.² The following set of parameters is provided within the system; these are named and outlined in §2.1.1–§2.1.6.

2.1.1 *Number of agents*

Steels (1997) reports on an experiment involving up to 50 agents, with ten active agents at any moment in time. There is an initial population, and gradual turnover not sensitive to communicative success (hence, this is a social model rather than a selectionist one). The experiments of Hurford (1989) involved a constant population size of 30 agents. Hurford (2000) supposes that at any one time there are four adults and one child language learner. The workbench presented here allows as an independent parameter the specification of an integer number of agents. Clearly, the capacity to experiment with a range of settings is desirable. Setting the value of the parameter at 1 essentially models the situation in which an agent evolves a language of thought in the Chomskyan sense, rather than an externally viable communicative tool.

Other agent-based features would be useful as parameters, given the sorts of experiments that have been conducted in the literature to date. Interaction constraints among agents are not modelled at present (for example, one might want to suppose that more than one agent can hear an utterance in a context at any one moment of time). Elimination of agents over time is not modeled.

²The system is implemented in *sicstus-Prolog*. It is available to other researchers to experiment with their own views on parameter settings, and for relevant enhancements that parallel our own development of the system.

2.1.2 Size of phoneme space

Human languages vary widely in the number of phonemes they make use of de Boer (1997) reports on a range of 11 to 141, with around 30 as a typical range for a human language. Clearly, there is an interaction between the number of discriminable sounds and the number of atomic signals available in a language, even if at a basic level, phonemes are meaningless, but give rise to the duality of patterning in which meaningful units may be formed (Carstairs-McCarthy, 1998a, 1998b). In the system described here, the total number of phonemes available to each agent may be specified, but constraints on potential co-occurrence of phonemes cannot be stipulated (cf. de Boer, 2000).

2.1.3 Semantic Space

One of the fundamental issues to be examined in simulating the emergence of natural language is in the estimating the space of possible meanings. Setting aside abstract entities and intensions, it is important when considering a simulation to know whether it assumes an open or closed semantic space in extensional terms: the number of objects that exist and the number of relations that they can stand in. Thus, there are two relevant dependant parameters: **Size of entity pool** and the **Size of relation pool**.

Note that the simulations of Hurford (1989) operate five objects and seven signals. Steels (1997) assumes 10 and Kirby (2000) assumes 100 possible meanings. While the results that they obtain are interesting, it seems that a richer semantic space is prerequisite to language: with a finite atomic space of meanings, compositional signals are unnecessary. Thus, in the system described here, any natural number of discriminable entity types may be named, and any number of instantiations of seven basic event types are possible.

Relation types are as follows: R^1 (one animate argument); R^2 (two animate args); R^3 (one human arg, one unconstrained arg, one animate arg); R^4 (two animate args and one inanimate); R^5 (one human arg and one relation arg); R^6 (one unconstrained arg); R^7 (one human arg, one unconstrained, one relation). Thus, given the number specified as an input parameter, a corresponding number of events types are constructed by distributing the number over these seven sorts of event types. A third dependent parameter, **Zipfian vs. random distribution of events**, determines whether the sorts of event are randomly distributed, or exhibit a greater systematicity. That is, in a Zipfian world certain events happen a great deal more often than other events, with a Zipfian distribution to them. Such a world might appear more structured than a randomly distributed world, and that underlying structure may lead to correspondingly more structure to utterances about the world. Nonetheless, it is a parameter to explore in its interaction with other parameters of the system. That the meaning space is open ended is ensured by event types which conform to schema (R^5) and (R^7) which are recursively structured (e.g. like believing one is seeing someone see some other event.) If one were to allow fourteen event types, and if a random distribution resulted in an even distribution of those

fourteen among the seven schema, then there would be four basic sorts of one place events that could happen, two constrained to have animate objects and two allowing inanimate arguments as well. Similarly, four event types would allow recursion. The discriminability of an event type is distinct from its occurrence. Thus, any instance of a recursive event can involve arbitrarily much recursive event structure. Hence, the space of possible things to communicate about when events of the constructed types start to happen is unlimited: the meaning space is unbounded, even though extensional.

2.1.4 Number and length of epochs

The **number of epochs** is one of a number of dependent parameters which jointly determine the duration of a simulation. **Epoch length** can be thought of as a span of communication among agents in which events happen, one agent produces a comment on the event and another agent interprets that comment relative to its own perception of the event. Given that this is about the association of signals to meanings for the very first time, it is reasonable to imagine that one parameter should be a sort of **forgetting threshold**, such that if a symbol-meaning association is not frequent enough during a stretch of conversation, then it is forgotten. Thus, an **epoch length** as set here is the point at which new, under-utilized, associations are forgotten. The forgetting threshold can be seen as mimicking the effect of lexical self-propagation among groups of agents and the random turnover of agents provided in Oliphant's (1999) social model. Section 3 outlines results of simulations in which the parameter is set to 0, thus providing a purely social/cultural model. The **number of epochs** is then the total number of such conversations to model.

2.1.5 Communicative Success

It is implicit in the discussion so far that some research in the area pursues a genetic model of language evolution (e.g. Hurford, 1989) and others favor a social model (e.g. Kirby, 2000). Both of these instances of those approaches, and, in fact rather a lot of simulations in the literature, assume a semantic transparency in which during learning stages the simulated interlocutors have access to each other's signal and meaning pairs. The rest of the system dynamics bootstrap from telepathy. Separate from whether interlocutors can know that they are 'talking about the same thing', it is possible to model success of communication. The least assumption-laden approach to the onset of communication is apparently fairly pessimistic: events happen, but agents individuate events according to their own perspectives, and comment on them accordingly. Because this is about the ontogenesis of language there is no good reason to assume that there are shared ways of mapping symbols onto meanings (unless one makes the assumption that in fact, everyone share's one's own mappings). In such a pessimistic setting, interlocutors can have a measure of success if they see the world the same way, but describe it differently. Language may have emerged to satisfy the needs of gossip more than survival (cf. Power,

2000), but success in actual communication probably does matter some times. The examples provided by Cushing (1994) about miscommunications involving aircraft crashes are striking in part because such spectacular ramifications of miscommunications so rarely ensue. Thus, in the simulation platform discussed here, it is possible to set at a value between zero and one the **Probability that success matters**, and within that, if any particular coin toss determines that success does matter, **Level of necessary success** necessary.

2.1.6 *Sensitivity to recency vs. frequency*

A final parameter that we mention here relates to how an agent constructs an utterance in response to an event or how an agent comes to interpret an utterance. The system assumes that creative language use is also a possibility, but that there is a tendency for each agent to explore its memory of past associations between symbols and meanings during production and interpretation. Because of the preponderance of recency and frequency effects in human behavior, it seemed reasonable to include agent sensitivity to one or other of these as a parameter in exploring differential impact on emergent system dynamics. An agent who is sensitive to recency rather than frequency has access only to local coordination processes as only the most recent prior event can have an impact on the current one, while frequency based agents have additional access to global co-ordination (cf. Garrod & Doherty, 1994)

Of these parameters, phoneme space and semantic discrimination capacity are explored in detail in §3. In addition to issues named above in which the system does not provide a mechanism for providing relevant constraints on parameters, other parameters are under exploration.

2.2 Measurable Quantities

Before providing details about the representations and key algorithms involved, we mention some of the system behaviors that are interesting to measured as a way of assessing the state of convergence on a communicative system: Level of understanding achieved in a communication; Average understanding of last 10 utterances; Homonymy ratio; Synonymy ratio; utterance complexity.

While the interlocutors themselves cannot inspect each others' meanings, it is possible to externally monitor their degree of communicative success. Even with pessimistic assumptions about the amount of miscommunication that might persist in the world, actual communication can be successful, unbeknownst to the interlocutors. Convergence on use of terms can emerge even if individuals assume only that they share meanings without attending to evidence either way. Success in communication is computed as an average over all communications and in the ten most recent utterances.

However, other meaning-based assessments of linguistic innovation are also important. Evidently, natural language tolerates more homonymy than synonymy; 'synonymy' here refers to a state in which one meaning can be denoted by a larger number of *basic* expressions. Natural language supports a larger

amount of synonymy at the phrasal level than lexical. If one were rationally designing a language one would prefer the alternative of more synonymy than homonymy to increase the likelihood of arriving at the same meaning from distinct signals rather than having the same signal yield many meanings. While we accept a point emphasized by Steels and Kaplan (1999) to the effect that homonymy is most acceptable when the intended meanings of homonym sets is maximally distinct, we nonetheless take it as an independent validation of linguistic properties within a simulation system for the number of synonyms to be smaller than the number of homonyms.

Any number of important measures of the system dynamics can be made, and here we have commented on measures we have been using to guide our understanding of the mutual impact of various parameter settings. Some of these measures could be calculated in different ways, and additional measures included.

2.3 Implementation Issues

The system does not assume that there is transparency of meaning, does assume that agents have the same space of possible meanings available to them, and the same range of basic concepts. However, agents need not segment events into constituent events in the same way, even though they decompose complex events into the same set of conceptual atoms. Entities, entity types and event types are represented as Prolog atoms, their numbers are determined by settings to input parameters. Events themselves are modeled as instances of Prolog lists pairing event types with arguments. Entities are classified as either human, animal or inanimate, with event types selecting arguments of type: human, animal, animate, inanimate, event type or without constraint.

Based on the numbers given as parameters, a corresponding number of event types and entities is created, assigning them to the possible categories as discussed in §2.1.3. Random distinct Prolog atoms are constructed for each, varying in form for human readability. An example event is this:

(1) `[ihdixos, spmg, davr, fizg];`

where `ihdixos` was an atom constructed to correspond to an event type with, it happens, an animate first argument, an inanimate second argument and an animate third argument. Animacy of the arguments actual arguments is randomly assigned, but relations are constructed to have expectations of which arguments can participate in its limited set of roles. Recall that some event types embed relations as arguments. Thus, although there are a finite number of events as provided by an input parameter, if one or more of those event types is of a sort that allows embedding, then during the course of simulation in which events happen and speakers comment on them, individual events may be arbitrarily complex. In the limit there is an infinite number of possible events even given finite specification of the number of event types. This is richer than boundless iterability of a finite set of events.

Perhaps more important is that the events do not themselves prejudice learning towards the context free grammar implicit in a predicate argument representation of an event with the functor providing the left-hand nonterminal and the arguments supplying the right-hand side of a production. When an event happens (e.g. (2)) the speaker selects a random element of the set of all sublists of the event, where each list of sublists corresponds to a perspective on the event. For example, (3) and (4) are two distinct perspectives an agent might have with respect to the example event.

- (2) [ihdixos, spmg, davr, fizg]
 (3) [[ihdixos], [spmg], [davr], [fizg]]
 (4) [[ihdixos], [spmg, davr], [fizg]]

The speaker and hearer are not constrained to have the same perspective.

Phonemes are also arbitrary atoms. A speaker, upon witnessing some event, uses pairs of phonemes to refer to each element of the event. At the outset, this involves invented pairs for the association, but over time, past experience interacts. Thus, one element of constraining structure imposed on the model, is that speakers try to talk about the entirety of their perspective on an event placed before them. However, through iteration between being a speaker and a hearer, a speaker may come to prefer a more complicated way of referring:

- (5) Event: [jilufks, furk]
 Uttered: [[[jilufks], [[r, q], [r, v]]], [[furk], [[e, y]]]]
 Heard: [[[vmpk], [[r, q]]], [[furk], [[r, v], [e, y]]]]

Here an event occurred, and a speaker used a four-phoneme complex,

- (6) [[r, q], [r, v]]

to pick out the event type and a two phoneme complex to pick out the argument. Thus, words are modeled as lists of pairs of phonemes unrestricted in length. There is sufficient space for duality of patterning to emerge without being built-in. No model of articulatory constraints on possible phoneme pairings is made.

The hearer has unobscured access to the same event commented on by the speaker. This models the sharing of cognitive possibilities among communicating agents. However, perspectival divergence is also possible. The hearer can partition the event differently. The onus on the hearer is not to find a phoneme sequence for each part of the perspective on the event, but something in the conceptual space for each phoneme sequence to mean. Another assumption that is not parameterized is that there is no noise—while the hearer may segment the signal differently (as above), the hearer has perfect access to the stream of phonemes uttered, but not segmentation or denotation.

The asymmetry in responsibility between speaker and hearer does not contradict the Saussurean perspective. Speaker is required to comment on everything in its perspective on the event; hearer is required to ground every word it segments from the signal either in its perspective on the shared event or on

past interpretations. Perspective is known to confound theoretical assumptions. Schober (1993) demonstrates the lack of speaker attention to hearer perspective in identifying use conditions for definites. The asymmetry we suggest here contradicts the Saussurean perspective only to the extent that those empirical findings do.

There is no initial grammar that constrains the system, neither explicitly, nor implicitly in the structure of semantic representations of events. There are only associations of meaning sequences with phoneme sequences.

2.3.1 Algorithms

The basic architecture of the system involves iterating through the following process based on the input parameters.³

1. Initialize:
 - clear memory, etc.
 - generate enough agents ^{π} , phonemes ^{π} , entities ^{π} and relations ^{π}
2. If out of Epochs ^{π} , show statistics & quit.
3. Forget any symbol/meaning pair from the last epoch that occurred no more often than the forgetting threshold ^{π}
4. Run another epoch.
 - (a) Some relation obtains at random
 - (b) A speaker comments on that event, note taken
 - (c) A hearer observes the event and interprets the utterance, note taken
 - (d) Omniscients observe the degree of common understanding
 - (e) If it's a situation falling under the probability that success matters ^{π} , and if success is less than a threshold percentage ^{π} ,
 - then ignore the utterance
 - else, note the symbol/meaning association (updating frequencies), and go to 4a
5. Goto 2

This algorithm involves iterations of speakers commenting on events, and hearers interpreting utterances in the context of a jointly witnessed event. These subroutines are outlined below. Both speakers and hearers have the capacity to innovate. Both speakers and hearers are influenced by the history of communication (modulo that which is forgotten because it happened so infrequently). That influence of history may be frequency based (depending on all prior utterances) or recency based (depending on the most recent utterance). Because over time speakers and hearers will interchange roles and be equally influenced by their own past interpretations as by associations put forward through their utterances, there is a communal knowledge base of past interpretations. It is important to point out that in our idealization there is

³A superscript is supplied for each quantity that is parameterized.

only one knowledge base of past utterances, modeling the agents' dual roles as speakers on some occasions and hearers on others.⁴ This is not a model of differential memories among the participants, nor is there a real model of information state beyond the immediate event. A richer extension of the system would include participant-indexed knowledge bases and information states, with updates, downdates, queries, disputes and acknowledgements.

Speaker comment on events

1. The event is modeled as the list of atoms given by a relation name and its arguments.

2. The speaker individuates that event.

This is modeled by:

- (a) This list of possible partitions of the event list is formed.
- (b) A random partitioning is selected from that list.

3. The speaker associates a symbol/phoneme sequence with the structured perceived meaning.

This is in partial relation to what has been uttered and construed before.

For each meaning segment in the speaker's individuated meaning:

- (a) Identify all the phonemes associated with it in the past
- (b) If that list is nonempty:
 - i. Choose a random element of it.⁵
 - ii. If some chance event happens, whose probability diminishes with the number of words uttered (and remembered) so far, then just invent a new phoneme sequence to associate with the meaning.⁶ Otherwise, the uttered phoneme sequence is the random choice among past events.

Otherwise, that meaning hasn't been noted before, so make up a new phoneme pair to associate with the meaning unit.

Hearer interpretation of an event

1. The interpreter hears an event and segments the stream into units.

2. Interpretation associates meanings with the segments in relation to the event that occurred and the history of past associations

3. The interpreter is either frequency ^{π} or recency ^{π} sensitive.

⁴One could well argue that this idealization is at best of an individual communicating with itself. To the extent that it's a valid argument, one must also agree that communicating with oneself is a nontrivial ability to have emerged. Recall, for example, Jaynes's (1976) hypothesis that the origin of human consciousness is in the integration of the hemispheres, the end of interhemisphere intraindividual discrete communication; also recall the arguments of Fodor (1983) and Chomsky's minimalism that the role of language is not so much interpersonal communication as interfacing to the rest of cognitive architecture.

⁵This implements frequency effects: it is not the most frequent pairing that is necessarily selected, but frequent pairings have higher probability of selection from memory than infrequent pairings.

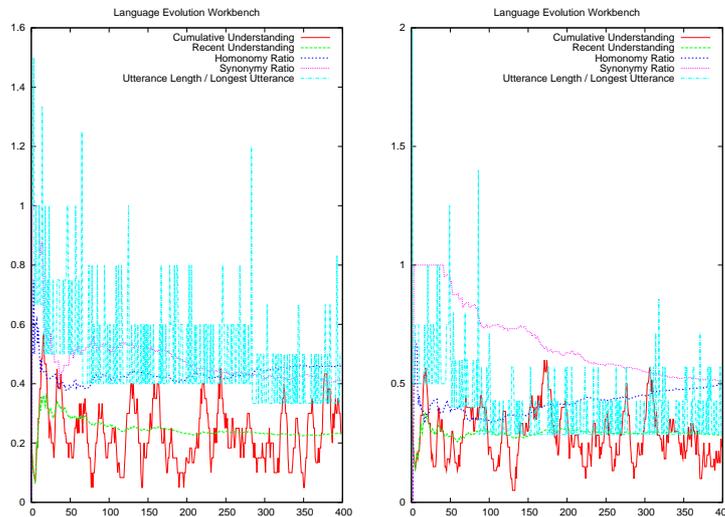
⁶Thus, innovation is possible at any stage, whether influenced by frequency ^{π} or recency ^{π} .

- (a) If there are no more “words” in the segmentation, association is done.
- (b) Find all past associations of meanings with the next word.
 - i. If frequency sensitive, select a random element of this list as the current interpretation of the word.
 - ii. Otherwise, if recency sensitive, select the last interpretation of the word as the current.
 - iii. If no past associations exist for the word, then on the basis of the hearer’s own individuation of the event, associate a meaning-segment with the speech segment.
- (c) Goto 3a

Notice that the model assumes that speakers have more opportunities for innovation than hearers, but that both have opportunity to innovate.

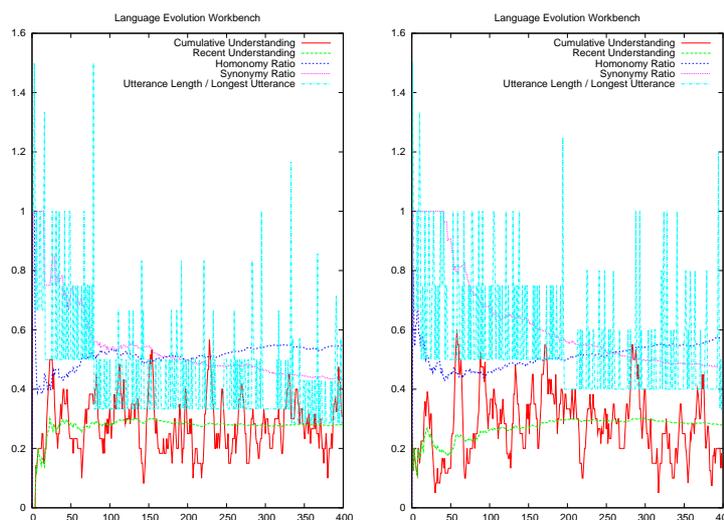
3 Sample Experiments & Results

Assume the following parameter settings: Number of agents: 100; Size of phoneme space: 41; Size of entity pool: 1000; Size of relation pool: 100; Number of epochs: 30; Attention span: 30; Forgetting threshold: 0; Probability that success matters: 25%; Level of necessary success: 75%; Sensitivity to frequency; Zipfian event distribution. A typical run of the simulation will have the behavior as graphed below on the left. The plot on the right holds all of the other parameters constant, but uses a random distribution of events types.



The two graphs both plot utterances on the x axis and an interpretation of the y axis as a percentage. The value exceeds 1 when an utterance exceeds the length of the longest utterance so far. In both cases, synonym and homonym ratios appear to cross path. In both cases, average recent understanding, the amount

of utterance signal and meaning pairs speakers and hearers agree upon, is relatively quickly in the 30% range, yet slightly better in the non-zipfian world. Note that these are not particularly demanding constraints on the likelihood that mutual understanding is important, or its level of importance when it is important. The following two are representative runs just like the above pair, but with understanding set at the inverse level: 75% chance that an utterance’s meaning matters, but when it does, a 25% level of sharing is necessary before the interlocutors move on to the next event.



These outputs are intended to indicate the sorts of experiments that one might conduct using the simulation platform. Additionally, it would be possible to plot other relevant values, perhaps localizing the evaluations to individuals rather than summing over all of them as in the plots depicted.

4 General Discussion

In some other theoretical work and also some simulations it is assumed that language evolution begins with perfect mappings between primal utterances and associated meanings, those mappings available to the language learner from the start of learning, and even as language itself begins its evolution. We do not feel that it solves any problems to assume that initial communications are perfect (which is what one has if both utterance and meaning are available to both speaker and hearer) and that what evolves are complications in syntactic features. In our model, this is a parameter: although we can observe the degree to which understanding has taken place, it isn’t necessarily the case that the interlocutors know both what was said and what was meant. A question we explore is whether under such pessimistic assumptions about interintelligibility a systematic language can emerge with actually successful communication, success rated by external observers.

We have presented an implemented system which has been used to explore parameters in simulations of linguistic evolution which might serve among the preconditions for the first appearance of language. We have shown here the perhaps surprising result that external observers of communicators who do not know if they are communicating successfully, under a range of reasonable parameter settings, tend to observe better than chance success in linguistic communication between the agents. We also observe that convergence on successful communication under reasonable parameter settings tends to happen very rapidly, within the first score of utterances. Unlike natural languages, and rather like one would expect from a rational design for an artificial language, the system tends to minimize homonymy and maximize synonymy. We present the system which is oriented to a social rather than a genetic model of linguistic evolution. The role of this paper is not so much to argue from any other particular simulations that any certain parameter settings are important, as it is to offer this system forward for others in the community to develop further and cross-validate research hypotheses. Woods (2003) demonstrates a number of experiments varying those parameters. We have since added questions and reflection (metalinguage).⁷

References

- Carstairs-McCarthy (1998a). Synonymy avoidance, phonology and the origin of syntax. In Hurford, J. R., Studdert-Kennedy, M., & Knight, C. (Eds.), *Approaches to the evolution of language. Social and cognitive bases*, chap. 17, pp. 279–296. Cambridge University Press.
- Carstairs-McCarthy, A. (1998b). *The Origins of Complex Language: An Inquiry into the Evolutionary Beginnings of Sentences, Syllables and Truth*. Oxford: Oxford University Press.
- Cushing, S. (1994). *Fatal Words: Communication Clashes and Aircraft Crashes*. University of Chicago Press.
- de Boer, B. (1997). Self organisation in vowel systems through imitation. In Coleman, J. (Ed.), *Computational Phonology, Third Meeting of the ACL SIGPHON*, pp. 19–25.
- de Boer, B. (2000). Emergence of sound systems through self-organization. In Knight, C., Studdert-Kennedy, M., & Hurford, J. (Eds.), *The Evolutionary Emergence of Language: Social Function and the Origins of Linguistic Form*, pp. 177–198. Cambridge University Press.
- Fodor, J. A. (1983). *The Modularity of Mind*. Cambridge: MIT Press.
- Garrod, S. & Doherty, G. (1994). Conversation, co-ordination and convention: an empirical investigation of how groups establish linguistic conventions. *Cognition*, 53, 181–215.

⁷Lucinda Longmore, MSc dissertation in preparation, TCD.

- Hurford, J. (1989). Biological evolution of the saussurean sign as a component of the language acquisition device. *Lingua*, 77, 187–222.
- Hurford, J. (2000). Social transmission favours linguistic generalisation. In Knight, C., Studdert-Kennedy, M., & Hurford, J. (Eds.), *The Evolutionary Emergence of Language: Social Function and the Origins of Linguistic Form*, pp. 324–352. Cambridge University Press.
- Jaynes, J. (1976). *The Origins of Consciousness in the Breakdown of the Bicameral Mind*. Princeton.
- Kirby, S. (1998). Syntax without natural selection: how compositionality emerges from vocabulary in a population of learners. In Knight (Ed.), *Second International Conference on the Evolution of Language*. Oxford: Oxford University Press.
- Kirby, S. (2000). Syntax without natural selection: how compositionality emerges from vocabulary in a population of learners. In Knight, C., Studdert-Kennedy, M., & Hurford, J. (Eds.), *The Evolutionary Emergence of Language*, pp. 303–323. Cambridge University Press.
- Oliphant, M. (1999). The learning barrier: moving from innate to learned systems of communication. *Adaptive Behavior*, 7(3-4), 371–384.
- Power, C. (2000). Secret language use at female initiation: bounding gossiping communities. In Knight, C., Hurford, J. R., & Studdert-Kennedy, M. (Eds.), *The Evolutionary Emergence of Language*, chap. 5, pp. 80–97. Cambridge University Press.
- Schober, M. F. (1993). Spatial perspective-taking in conversation. *Cognition*, 47, 1–24.
- Smith, K. (2002). Natural selection and cultural selection in the evolution of communication. *Adaptive Behavior*, 10(1), 25–44.
- Steels, L. & Kaplan, F. (1999). Bootstrapping grounded word semantics. In Briscoe, T. (Ed.), *Linguistic evolution through language acquisition: formal and computational models*. Cambridge University Press.
- Steels, L. (1997). Synthesizing the origins of language and meaning using co-evolution, self-organization and level formation. In Hurford, J. (Ed.), *Evolution of Human Language*. Edinburgh University Press.
- Vogel, C. & Woods, J. (2002). Simulation of evolving linguistic communication among fallible communicators. In Hurford, J. & Fitch, T. (Eds.), *Proceedings of the Fourth International Conference on the Evolution of Language*, p. 116. Harvard University, Cambridge, MA.
- Woods, J. (2003). Declaratives, interrogatives, semantic space and the emergence of communication. Master's thesis, Computational Linguistics Lab, Trinity College, University of Dublin.