Modelling Relevance-Driven Language Evolution

Stefan Hoefler

Institute of Computational Linguistics, University of Zurich
Binzmühlestrasse 14, 8050 Zurich, Switzerland
hoefler@cl.uzh.ch

Abstract. Computational modelling has proven a useful method to study the emergence of language-like communication systems. However, most existing models abstract away from the two facts that (1) language use exhibits pragmatic plasticity[1] and (2) linguistic knowledge is an integral part of human conceptual knowledge. This paper introduces a basic architecture for a model that overcomes these shortcomings by incorporating elements of Relevance Theory and Cognitive Semantics.

Key words: simulating language evolution, cognitive linguistics

1 Introduction

In the obvious absence of primary evidence, the study of the mechanisms of the evolution of human language has, to a large part, been conducted by means of computer simulations (see [Cangelosi and Parisi, 2002] and [Briscoe, 2002] for overviews). In such simulations, computational agents iteratively engage in communicative interactions. An agent possesses (1) some sort of knowledge base, (2) an algorithm to produce and interpret signals, and (3) a learning algorithm to update his knowledge base on the basis of experienced communicative exchanges. In each iteration of a simulation, agents are given a randomly generated meaning to express and, upon doing so (or upon observing another agent doing so), adapt their knowledge base accordingly. Such simulations are used to study how certain features of human language (e.g. recursion, compositional syntax) emerge in the agents' knowledge base over time.

Existing computational models commonly limit the agent's knowledge base to linguistic knowledge, which is typically represented as a code in the form of a generative grammar. The acts of producing and interpreting a signal are thus reduced to mere processes of encoding and decoding. The employed learning

---

[1] It was when Michael Hess introduced me to computational semantics that I first became aware of this fact. The insights I gained from his comprehensive and well-structured lectures later inspired me to develop computational models to study the phenomenon's role in language evolution and to pursue a doctoral dissertation in computational evolutionary linguistics. The present paper is an offshoot of this work.

algorithms are consequently limited to grammar induction: new generations of agents added to a simulation during runtime induce their own generative grammars on the basis of the data produced by earlier generations of agents. While this class of models has yielded valuable insights about possible mechanisms of language evolution—e.g. how recursive syntax [Kirby, 2002] or syntactic ambiguity [Hoefler, 2006] can emerge from iterated grammar induction—they suffer from two shortcomings.

For one, they abstract away from the fact that human utterance production and interpretation involve inference from context. Signal meanings (the meanings that a grammar or linguistic code conventionally associates with a signal) hardly ever fully specify the speaker meaning (the meaning that is actually communicated in a concrete situation). Linguistic communication rather exhibits pragmatic plasticity [Hoefler, 2009]: signal meanings usually underspecify and/or overspecify the speaker meaning, i.e. they are less or more specific than the meaning the speaker actually conveys. A hearer will infer additional relevant information from the context and ignore irrelevant aspects of the "literal" meaning of a signal on the basis of what she recognises as knowledge shared with the speaker [Clark, 1996]. An essential part of such common ground is a shared understanding of what the goal of the communicative interaction is, which in turn, leads to shared assumptions about what constitutes relevant information in a given situation [Grice, 1975].

A second shortcoming of existing models of language evolution is that they isolate linguistic knowledge. Cognition-based studies of human language, however, show that linguistic knowledge is not separate from non-linguistic knowledge but rather constitutes an integral part of broader human conceptual knowledge [Evans and Green, 2006]. Furthermore, it has been suggested that just like lexical items, syntactic constructions too can and should be understood as symbolic associations between a form (concrete in the case of lexical items and schematic in the case of syntactic constructions) and a meaning [Goldberg, 1995, Croft, 2001]. Both linguistic as well as non-linguistic knowledge can thus be represented as an inventory of associations between concepts [Langacker, 2008].

In this short paper, I will sketch a general architecture for a model of language evolution that overcomes the aforementioned shortcomings. To this aim, I will first introduce a prototypical implementation of an agent's knowledge base that integrates linguistic and non-linguistic knowledge and describe how relevance-driven utterance production and interpretation can be modelled. Then, I will briefly outline how such a model can be used to study the emergence of syntactic constructions in the course of language evolution.

2 A Relevance-Driven Model

The model to be introduced represents an agent's knowledge base as an inventory of conceptual units. A conceptual unit is an entrenched psychological structure that an individual can access as a whole (as a gestalt) in a largely automatised fashion, without any constructive effort being necessary [Langacker, 1987,

p. 57f.]. The model's agents store conceptual units together with their degree of entrenchment, i.e. together with a measure that indicates the ease with which the unit can be accessed:[2]

```
unit([elephant(_)], 30).
```

Psychological studies show that the entrenchment of a conceptual unit increases through use and decreases through lack of use [Croft, 2000]. In the model, the entrenchment value of a unit at a time t is therefore calculated as the number of times that unit has been used at time t, divided by the number of consecutive iterations in which it has not been used anymore [Hoefler, 2009, p. 160].

If the agent observes units occurring together, he will combine them to form a new, complex unit that is added to his knowledge base—with an (as yet) minimal entrenchment value. However, "[i]t is important to observe that when a complex structure coalesces into a unit, its sub-parts do not thereby cease to exist or be identifiable as substructures [...]" [Langacker, 1987, p. 59]:

```
unit([scotland(_), rain(_)], 20).
unit([italy(_), rain(_)], 5).
```

One crucial property of the model is that conceptualisations of phonological entities (sound sequences, linear order, etc.) are not treated any differently from other concepts: as [Saussure, 1916] and [Langacker, 1987] have long pointed out, psychologically, both forms and meanings are conceptualisations of objects or states of affairs in the world. Consequently, an agent's knowledge base can come to contain associations between co-occurring non-phonological concepts (e.g. the above association between Scotland and rain) as well as associations between linguistic forms and the meanings with which they have been observed to co-occurr (e.g. the associations between */pen/* and its meanings 'inked writing utensil' and 'fenced enclosure'). In accordance with recent cognitive studies, the model thus fully integrates linguistic and non-linguistic conceptual knowledge.

```
unit([phon_pen(_), writing_utensil(X), inked(X)], 75).
unit([phon_pen(_), enclosure(X), fenced(X)], 12).
```

It needs to be noted that beyond the aforementioned features, the model does not constrain the experimenter to any specific set of concepts. It rather offers a general framework that can be used to investigate the impact of different hypothetical conceptual spaces on the evolution of language. The concepts used in the examples in this paper merely serve the purpose of illustrating the architecture of the model.

To perform inferences like those involved in language use, the agent needs to have access to additional, interaction-specific knowledge. In accordance with Relevance Theory [Sperber and Wilson, 1995], the model thus presupposes that in each communicative interaction, the interlocutors are capable of determining

---

[2] In accordance with the notation used in Cognitive Grammar [Langacker, 2008], conceptual units are marked with square brackets.

what would constitute relevant information if communicated. Relevance may be a matter of degree [Sperber and Wilson, 1995]. Imagine, for instance, that two people in a car approach some traffic lights, and the driver asks the passenger what colour the traffic lights are. In this situation, both `[red(_)]` and `[green(_)]` would constitute relevant information. However, the information that the traffic lights are red would be somewhat more relevant as it would necessitate some action on the part of the driver (stopping), while he could continue at the current speed if the lights were green. In the agent's knowledge base, such knowledge would be represented, for instance, as follows:

```
relevant([red(_)], 0.6).
relevant([green(_)], 0.4).
```

In its current state, the model does not specify how such knowledge is gained—doing so, e.g. by implementing the respective relevance-theoretic mechanisms [Sperber and Wilson, 1995], would constitute a project on its own—but presupposes that it is available in each communicative interaction. Its inference is simulated by providing the agent with a randomly generated set of facts of the type shown above at the beginning of each iteration of a simulation.

The algorithm that the agent employs for utterance interpretation and signal production largely corresponds to the descriptions of language use given by Relevance Theory [Sperber and Wilson, 1995] and Cognitive Semantics [Fauconnier and Turner, 2002]. These theories view linguistic signals as cues that do not carry meaning per se but merely serve as access points to the interlocutor's conceptual knowledge. An observed signal triggers the (neurological) activation of certain concepts, which may activate further concepts, and so on. The process stops once a conceptual structure has been activated that constitutes relevant information in the context of the given communicative interaction: "[c]heck interpretive hypotheses in order of their accessibility—that is, follow the path of least effort until an interpretation that satisfies the expectation of relevance is found; then stop" [Carston, 2004, p. 822].

The present model implements the activation of units as conceptual blending [Fauconnier and Turner, 2002]: one conceptual unit activates another conceptual unit by blending with it. If more than one unit can be blended with, the one with the highest degree of entrenchment (i.e. the one that can be accessed with least effort) is activated first. Two units can blend if they share at least one component, a so-called "anchor." Blending is realised as graph-unification relative to that anchor:

```
blend(Unit1, Unit2, Blend) :-
     member(Anchor, Unit1),
     member(Anchor, Unit2),
     union(Unit1, Unit2, Blend).
```

To illustrate the utterance interpretation algorithm, imagine e.g. that in the above example, the passenger answered the driver's question about the colour of the traffic lights by uttering "Robin." Let us further suppose that at the time the interlocutors shared the following knowledge base:

```
relevant([red(_)], 0.6).
relevant([green(_)], 0.4).
unit([phon_robin(_), robin(X), bird(X)], 60).
unit([phon_robin(_), robin(X), firstname(X, _)], 45).
unit([robin(X), bird(X), breast_of(Y, X), red(Y)], 78).
```

In this context, the observation of the signal `[phon_robin(_)]` would trigger the following chain of activation in the agent:

1. Does `[phon_robin(_)]` activate (i.e. blend with) a relevant meaning? No.
2. The next unit that is activated is `[phon_robin(_), robin(X), bird(X)]`. Does the blended structure `[phon_robin(_), robin(X), bird(X)]` activate a relevant meaning? No.[3]
3. The next unit that is activated is `[robin(X), bird(X), breast_of(Y, X), red(Y)]`. Does the blended structure `[phon_robin(_), robin(X), bird(X), breast_of(Y, X), red(Y)]` activate a relevant meaning? Yes, this time, the relevant meaning `[red(_)]` can be activated.
4. In the present context, `[red(_)]` is thus assumed to be the interpretation of the observed signal.
5. In the subsequent process of learning, the entrenchment of each conceptual unit involved in the interpretation of the observed signal is increased by one, and a new unit combining the observed signal and the communicated meaning is added to the agent's knowledge base:
   `unit([phon_robin(_), red(_)], 1).`

With the addition of this new unit, the signal */robin/* can now directly activate the concept RED in future usage events, without the detour via the bird.[4]

Since the components of a signal are signals themselves and each signal comes with the presumption of relevance [Sperber and Wilson, 1995], the interpretation of signals with more than one component is the first relevant meaning that can be reached from *each* of the individual components of the signal. Imagine, for instance, a situation where the hearer has the following knowledge base:

```
relevant([chicken(X), escape(E), agent(X, E)], 0.6).
relevant([chicken(X), sleep(E), agent(X, E)], 0.4).
unit([phon_chicken(_), chicken(_)], 93).
unit([phon_sleep(_), sleep(_)], 91).
```

Now imagine that the signal `[phon_chicken(_), phon_sleep(_)]` is produced in this context. The hearer simulated by our model will take the first component, `[phon_chicken(_)]`, and activate units until he is able to activate a relevant meaning, which, in this case, will be `[chicken(X), escape(E), agent(X, E)]`. Then, he will interpret the second component of the signal, `[phon_sleep(_)]`, which will lead him to a different relevant meaning, namely `[chicken(X),`

---

[3] Had the linguistic form */robin/* been used in its literal sense, the interpretation algorithm would have been able to activate a relevant meaning at this point.

[4] Note that this requires that the algorithm performs a breadth-first search.

`sleep(E), agent(X, E)]`. By backtracking, he will now try to find the next possible interpretation for the first component of the signal, which will also lead him to `[chicken(X), sleep(E), agent(X, E)]`. This meaning is the first relevant meaning that is accessible from each component of the signal and will thus, in the given context, be assumed to be the interpretation of that signal.

The examples show a crucial property of the proposed interpretation algorithm: the same mechanisms are employed no matter whether the produced signal is used in its literal sense, or metaphorically (as in the traffic lights example), or in an underspecified way (as in the chicken example). The algorithm is thus consistent with the claim that literal and figurative language use form a continuum [Langacker, 1987, Sperber and Wilson, 1995]. In [Hoefler and Smith, 2009], we have argued that such a unified model of language use can account for both the emergence of symbolism and the process of grammaticalisation.

The corresponding production algorithm is equivalent to the interpretation algorithm, with the exception that its input is not a signal but an intended meaning, and that units are activated not until a relevant meaning is found but until a producible concept has been reached. The production algorithm can be described as follows:

1. Take the first component of the intended meaning, follow its path of activation until a producible concept is found; stop and produce that concept as a signal.
2. Check if the produced signal really will be interpreted in the intended way by feeding it back in to your own interpretation device; if the resulting interpretation is identical to the intended meaning, stop.
3. Else, repeat steps 1 and 2 for the remaining components of the meaning until the predicted interpretation corresponds to the meaning you intend to convey.

The described production algorithm illustrates that the proposed model has the capacity to simulate different levels of theory of mind (which several recent studies [Tomasello, 2003] identify as a pivotal factor in the evolution of language). If the speaker does not entertain a theory of mind with regard to his interlocutors, the employed production algorithm only consist of step 1 above. If the speaker makes the assumption that his interlocutors' knowledge is identical to his own, the production algorithm also includes steps 2 and 3. Finally, we can model speakers who realise that their interlocutors' knowledge can differ from their own and who entertain hypotheses about what piece of knowledge is shared. To do so, we add to each unit the list of individuals with which that unit is shared. The production algorithm then only makes use of units that are recognised as being shared with the current interlocutor.

3 Towards the Emergence of Syntax

We now briefly turn to the question of how the introduced model can be applied to simulate the emergence of syntactic constructions. While classical generative approaches to language postulate a strict distinction between syntactic rules and lexicon, more recent, cognitive studies have found that the two actually must be assumed to form a continuum [Goldberg, 1995, Croft, 2001]. Syntactic constructions are viewed as symbolic units, i.e. form-meaning associations, with schematic forms. To explain how word order can become fixed in the course of language evolution, and how specific types of word order can come to take on specific meaning, we have to make the phonological concept of linear precedence explicit. The example below shows a signal with three components: the word */chicken/*, the word */sleep/*, and the fact that the two were uttered in linear order. If we assume that this signal was used to convey the information that some chicken `x1` sleeps, then the new unit shown below is added to the agent's knowledge base after the usage event.

```
Signal:    [phon_chicken(V), phon_sleep(W), lp(V, W)]
Meaning:   [chicken(X), sleep(E), agent(X, E)]
New unit:  unit([phon_chicken(V), phon_sleep(W), lp(V, W),
                 chicken(X), sleep(E), agent(X, E))], 1).
```

With this new unit added, the concept LINEAR PRECEDENCE can now be used to activate the concept AGENT in future usage events: like lexical material, this syntactic element of form can then be employed as a signal to trigger the inference of some specific meaning, as suggested by [Hopper, 1987]. Thus, syntax can gradually attain a role in conveying meaningful utterances. Future experiments will have to show what conceptual spaces are required for more complex syntactic patterns to emerge from the accumulation of complex conceptual units like the one shown above.

4 Conclusion

The aims of computer simulations are different from those of NLP applications. Trying to evaluate them for efficiency or other measurable properties would be beside the point; unlike NLP applications, they aim for psychological plausibility and explanatory power. This paper has introduced an architecture for a computational model of language evolution that incorporates two aspects of human language commonly ignored by existing models: (1) language use exhibits pragmatic plasticity, i.e. signals frequently under- and/or overspecify the communicated speaker meaning, and (2) linguistic knowledge is an integral part of general human conceptual knowledge. The devised relevance-driven algorithm for utterance interpretation reflects the cognitive underpinnings of linguistic communication more closely than the mechanisms of encoding and decoding used in most existing computer simulations. Moreover, a greater explanatory capacity is achieved: the model can simulate not only literal but also metaphorical and underspecified language use. Finally, the model has the potential to simulate the emergence of language-like structure from general human conceptualisations because it integrates linguistic and non-linguistic knowledge. It thus offers a platform that is well adapted for experiments to study what sort of conceptual spaces are required for language-like systems to emerge and evolve.

References

[Briscoe, 2002] Briscoe, T., editor (2002). *Linguistic Evolution through Language Acquisition: Formal and Computational Models.* Cambridge University Press, Cambridge.

[Cangelosi and Parisi, 2002] Cangelosi, A. and Parisi, D., editors (2002). *Simulating the Evolution of Language.* Springer, London.

[Carston, 2004] Carston, R. (2004). Explicature and semantics. In Davis, S. and Gillon, B., editors, *Semantics: A Reader*, pages 817–845. Oxford University Press, Oxford.

[Clark, 1996] Clark, H. (1996). *Using Language.* Cambridge University Press, Cambridge.

[Croft, 2000] Croft, W. (2000). *Explaining Language Change: An Evolutionary Approach.* Longman, Harlow.

[Croft, 2001] Croft, W. (2001). *Radical Construction Grammar: Syntactic Theory in Typological Perspective.* Oxford University Press, Oxford.

[Evans and Green, 2006] Evans, V. and Green, M. (2006). *Cognitive Linguistics: An Introduction.* Edinburgh University Press, Edinburgh.

[Fauconnier and Turner, 2002] Fauconnier, G. and Turner, M. (2002). *The Way We Think: Conceptual Blending and the Mind's Hidden Complexities.* Basic Books, New York.

[Goldberg, 1995] Goldberg, A. (1995). *Constructions: A Construction Grammar Approach to Argument Structure.* University of Chicago Press, Chicago.

[Grice, 1975] Grice, H. P. (1975). Logic and conversation. In Cole, P. and Morgan, J., editors, *Syntax and Semantics 3: Speech Acts.* Academic Press, New York.

[Hoefler, 2006] Hoefler, S. (2006). Why has ambiguous syntax emerged? In Cangelosi, A., Smith, A. D. M., and Smith, K., editors, *The Evolution of Language: Proceedings on the 6th International Conference of the Evolution of Language*, pages 123–130, Rome. World Scientific Press.

[Hoefler, 2009] Hoefler, S. (2009). *Modelling the Role of Pragmatic Plasticity in the Evolution of Linguistic Communication.* PhD thesis, The University of Edinburgh.

[Hoefler and Smith, 2009] Hoefler, S. and Smith, A. D. M. (2009). The pre-linguistic basis of grammaticalisation: A unified approach to metaphor and reanalysis. *Studies in Language*, 33(4):883–906.

[Hopper, 1987] Hopper, P. J. (1987). Emergent grammar. *Berkeley Linguistics Conference (BLS)*, 13:139–157.

[Kirby, 2002] Kirby, S. (2002). Learning, bottlenecks and the evolution of recursive syntax. In Briscoe, T., editor, *Linguistic Evolution through Language Acquisition: Formal and Computational Models*, chapter 6. Cambridge University Press, Cambridge.

[Langacker, 1987] Langacker, R. W. (1987). *Foundations of Cognitive Grammar: Theoretical Prerequisites*, volume 1. Stanford University Press, Standford, California.

[Langacker, 2008] Langacker, R. W. (2008). *Cogntive Grammar: A Basic Introduction.* Oxford University Press, Oxford.

[Saussure, 1916] Saussure, F. d. (1959/1916). *Course in General Linguistics.* Philosophical Library, New York.

[Sperber and Wilson, 1995] Sperber, D. and Wilson, D. (1995). *Relevance: Communication and Cognition.* Blackwell, Oxford, second edition.

[Tomasello, 2003] Tomasello, M. (2003). *Constructing a Language: A Usage-Based Theory of Language Acquisition.* Harvard University Press, Cambridge, Massachusetts.