Models of Language Evolution: Does the Math Add Up?
14 April 2010, Utrecht, the Netherlands.
Venue: Drift 21, Sweelinckzaal

A workshop at the 8th International Conference on the Evolution of Language, 14-17 April 2010.

Program

10h00 welcome & general introduction Bart de Boer & Willem Zuidema: *MODELS OF LANGUAGE EVOLUTION: DOES THE MATH ADD UP?*
10h20 Tao Gong (Linguistics, MPI/Leipzig): *A BRIEF REVIEW OF COMPUTATIONAL MODELING IN EVOLUTIONARY LINGUISTICS*
10h55 Joachim de Beule (Artificial Intelligence, ULB/Brussel): *FLUID CONSTRUCTION GRAMMAR AND ARTIFICIAL CHEMISTRY*
11h30 Martin Bachwerk & Carl Vogel (Computational Linguistics, Trinity College/Dublin): *MODELLING SOCIAL STRUCTURES AND HIERARCHIES IN LANGUAGE EVOLUTION*
12h05 coffee break
12h20 Garrett Mitchener (Mathematics, Charleston College/SC): *SIMULATING THE EVOLUTION OF COMBINATORIAL PHONOLOGY*
12h55 Michael Franke (Linguistics, U. of Tuebingen): *PERFORMATIVE MEANING & CLAUSE-TYPE EVOLUTION*
13h30 lunch break
14h30 Dan Dediu (MPI/Nijmegen): *BAYESIAN MODELS OF LANGUAGE EVOLUTION: THE NEW UNIFYING PARADIGM?*
15h05 Monojit Choudhury (Microsoft Research, Bangalore) & Animesh Mukherjee (ISI, Turin): *COMPUTATIONAL MODELING OF LANGUAGE EVOLUTION: GAPS AND CHALLENGES*
15h40 coffee break
15h55 Michael Arbib (Computer Science/Neuroscience, U. of Southern California/L.A.): *MODELING LANGUAGE EVOLUTION: A GENERAL FRAMEWORK WITH THE MIRROR SYSTEM HYPOTHESIS AS A SPECIAL CASE*
16h30 general discussion
16h55 round up

Utrecht City Center Map

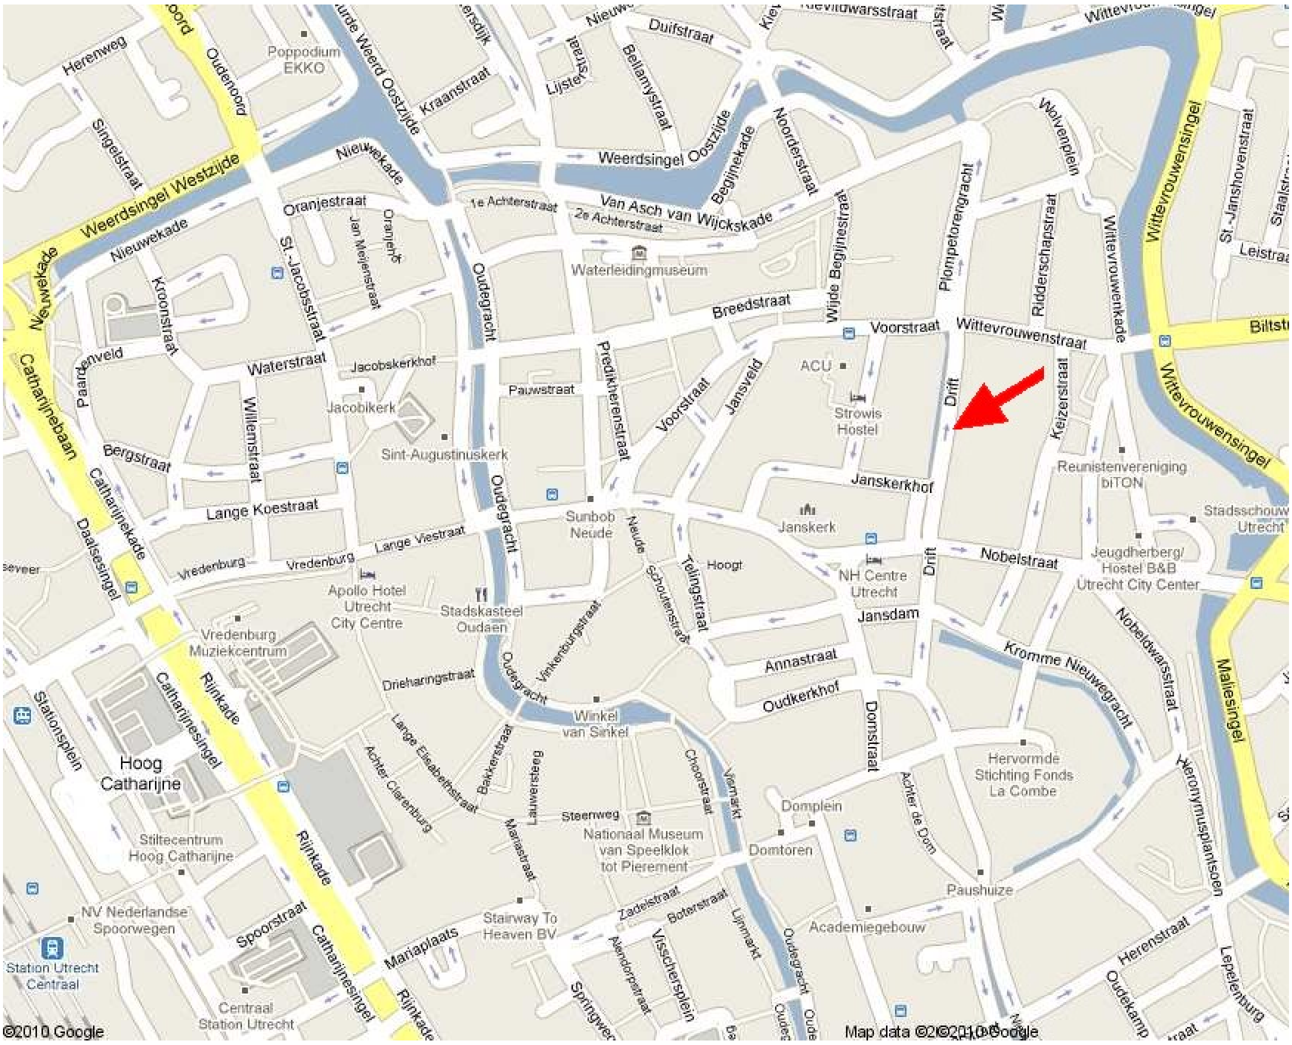

The workshop venue is 'Drift 21', room 'Sweelinckzaal', indicated on the map with the red arrow.

The other workshops and the opening ceremony take place in the 'Academiegebouw', indicated at the bottom right of the map. This is only a short walk away.

The central station is at the bottom left, also walking distance. Bus 11 to the university campus has a stop at the 'Janskerkhof', a few steps from the workshop venue.

Models of Language Evolution:

Does the Math Add Up?

Bart de Boer
Willem Zuidema
{b.g.deboer, w.h.zuidema}@uva.nl

1 — Introduction

The last two decades have seen an explosion of interest in mathematical and computational models of language evolution. Formal modelling is seen by increasingly many in the field as an approach that ensures internal consistency of evolutionary scenarios. However, there has been little attention to the question of how well the many different models fit together. Are they consistent with and complementary to each other? Is there a series of models that really covers the evolutionary emergence of modern language from a prelinguistic, ancestral state? Are the assumptions that go into a particular model, if not (yet) supported by empirical findings, made plausible by results from other models?

In this paper, we argue that these problems deserve much more attention than they currently receive. For sustaining the success of modelling approaches in language evolution research, it is crucial that models start living up to their promise: modellers must make explicit how their models fit in with other work in complete scenarios on the origins of language(s), and how their modelling results affect judgments of plausibility of one scenario against another. Moreover, they must do so based on careful consideration of other work, without overstating their results and misusing the prestige that comes with mathematical approaches.

Our arguments are based on a particular view on the role of modelling in scientific research in general, and in "historical" research fields with a paucity of direct evidence in particular. In section 2, we will therefore start with some considerations about the methodology of modelling in our field. To ground the discussion, however, we will quickly move to concrete examples. In section 3 we will discuss the contributions and shortcomings of models of the evolution of speech. In section 4 we will then draw some general lessons from this case study, and sketch an agenda for future research in the language evolution modelling field at large.

2 — Modelling methodology

From the great many distinctions one can make between different model studies, there are three particularly useful ones that also allow us to establish some common terminology and formulate the goals of this paper. The first is a distinction based on function, between predictive models and explanatory models (Gilbert & Troitzsch, 2005). Predictive models try to model a system as accurately as possible, and to make accurate predictions about the real system's behaviour, as in weather forecasts for example. Predictive models can also be used to reconstruct behaviour in the past, and could for example be used in reconstructing the spread of language families or of particular instances of language change (e.g., Landsbergen, 2009). Explanatory models, in contrast, aim to increase insight in a phenomenon. Explanatory models are generally much more abstract and further removed from reality than predictive models. The phenomenon under study is not modelled in all its detail, but instead only its essentials are modelled. Crucially, what counts as 'essential' very much depends on the research question, and simplifications that are appropriate for one question can be totally indefensible for another. Good explanatory models, moreover, explain the phenomenon of interest in terms of lower-level phenomena that can, at least in principle, be independently motivated (models that simply reproduce the phenomenon of interest without providing such an explanation are sometimes called *phenomenological models*).

The second important distinction is one based on form, between mathematical and computational models. The distinction is not always strict, but mathematical models tend to be the most abstract and to strip down phenomena to their barest essentials. Typically (but not exclusively), mathematical modelling papers provide both a formalization of a phenomenon (e.g., using matrix algebra, logic, differential equations) and proofs about properties of the formal system. Such proofs are, by definition, universally valid and allow inferences about specific cases (deduction), although the simplifications necessary to arrive at a proof often greatly limit the applicability.

Computational models tend to be much more concrete and complex. Phenomena are formalized in a programming language, and the resulting programs studied experimentally. From different runs with different parameter settings, the modeller tries to infer general properties of the formal system (induction). The programs can be very complex, allowing for models with fewer abstractions but often barring analytic proofs. In some cases, computational models are used to investigate versions of a mathematical model that are too complicated to study analytically (including *numerical models*, that are defined algebraically but studied using numerical methods on the computer).

A third major distinction concerns the validation of models: we distinguish between internal validation and external validation. Internal validation is about demonstrating that the phenomenon of interest indeed follows from the stated assumptions, and mathematical proof provides its most powerful form. This is much harder to achieve with computer models, although extensive testing and systematic exploration of the parameter space of a computational model can lead to a great degree of confidence. External validation is about checking whether the stated and unstated assumptions are supported by empirical evidence, or by the outcome of other, independent models, and whether the model's predictions are confirmed in the real world. As computational models are often formulated in more concrete terms, it tends to be easier to achieve external validation.

In this paper, we are mainly concerned with the external validation of explanatory models, which in all cases requires an interpretative step: explanatory models have, by definition, abstracted away many details of the phenomenon of interests, making it a matter of judgment whether abstractly formulated assumptions and prediction are supported by concrete evidence. In historical research fields such as the origins of language (and likewise, the origins of life or the universe), external validation is further complicated by the fact that there is little direct evidence about which assumptions and predictions are valid. External validation is thus only achievable by *model sequencing*: assumptions and prediction of any particular model are validated mainly by results from other models, and only at various points in a string of models do empirical results come into play.

Moreover, because this field deals with complicated phenomenon for which the appropriate simplifications haven't been established yet, modelling research should employ *model parallelisation*: for any particular phenomenon, researchers should develop multiple formalisations, compare results and relate observed differences to explicit and implicit assumptions embodied in these alternative models. Hence, modellers in language origins research must – much more than is currently practiced – work out relations between different models, whether they stand in sequence or in parallel to each other.

These observations may seem straightforward to many readers, but we find that much modelling work suffers from lack of clarity on these conceptual foundations. In the following section we will further explore our methodological points in the context of models of the evolution of speech, which, whilst home to some great controversies, provides a good case study as it has an abundance of easily interpretable data and a strong modelling tradition.

3 — The evolution of speech and repertoires of phonemes

When studying the evolution of speech in relation to the evolution of language, the focus is usually on the differences between modern humans and the hypothetical latest common ancestor (henceforth, LCA) of humans, chimpanzees and bonobos. The vocal abilities of the LCA are inferred from the abilities that humans, chimpanzees and other apes share or do not share. From such comparisons, it can be derived that the LCA had a repertoire of calls for communicative purposes, and therefore a limited ability to modulate the vocal tract. However, it most likely had a vocal anatomy more comparable to that of chimpanzees and vocal cords comparable to those of chimpanzees and gorillas. The LCA did not, it seems, have modern human's descended larynx, it had less voluntary control over breathing (MacLarnon & Hewitt, 1999) and probably did have supralaryngeal air sacs. As all modern apes only have limited voluntary control over their vocalizations, only learn their vocalizations to a very limited extent and lack internal (combinatorial) structure in their calls, it can be assumed that this was also the case for the LCA.

Modern humans, on the other hand, do have a descended larynx, have voluntary control over speech (but much less so over emotional utterances), and have a large learned repertoire of linguistic utterances. Moreover, those utterances have complex internal structure that is used productively, and there are regularities in the repertoires of speech sounds that humans use (the phonological universals). The challenge for research of the evolution of speech is to give an account of how the modern phenotype evolved from the LCA's phenotype: i.e., how did the descended larynx, voluntary control, vocal learning, combinatorial phonology and phonological universals evolve? A key issue here is to what extent the evolutionary changes should be considered adaptations for language, or to what extent they evolved for other reasons.

Computer models (and to some extent mathematical models) have been used for a long time to investigate such issues – but in the existing literature (as reviewed in de Boer, 2005; de Boer & Fitch, *in press*) there are some striking gaps in the range of topics considered and some disturbing confusions about the role of various models. The most studied topics are the evolution of the vocal tract (Lieberman & Crelin, 1971; Boë *et al.*, 2002; de Boer, 2009) and the emergence of phonological universals (de Boer, 2000b; Oudeyer, 2005; Zuidema & de Boer, 2009); the evolution of voluntary control, vocal learning and combinatoriality have received much less attention in the modelling literature, and the issue of how models of these different aspects fit together has been almost completely ignored.

Starting point for many models of how speech evolved are models of how speech perception and production works in human adults. Surveying the literature, we quickly find that many models that have been developed for the study of human speech are not necessarily directly usable in the study of the evolution of speech. Illustrative examples from modelling the acoustic production of speech are the 3-parameter model (Stevens & House, 1955; Fant, 1960), the couple mass-spring model (Dudgeon, 1970; Ishizaka & Flanagan, 1972) of the vocal cords and the source-filter model of speech production (Fant, 1960). These are simplified, explanatory models of the human vocal tract, the human vocal cords and the (lack of) interaction between the human vocal cords and the vocal tract, respectively.

These models are well established in phonetics, and provide valuable insights in the process of speech production. However, researchers in the evolution of speech cannot simply reuse these models to represent properties of vocal tracts of our evolutionary ancestors or of other species (see the discussion about Riede *et al.*, 2005 in Lieberman, 2006); doing so is misunderstanding the explanatory nature of the existing models, that involved simplifications which were very helpful for understanding speech production but are specific to human adult vocal tracts. It is, in fact, unlikely that ape-like vocal tracts can make the deformations of the vocal tract that are assumed by the 4-tube model, and it is clear that the acoustic effects of supralaryngeal air sacs are not captured by it. It is further unknown whether chimpanzee-like vocal cords work in the same way as human vocal cords, and whether in chimpanzee-like vocalizations the vocal cords can really be considered acoustically independent of the vocal tract. Simplifications made in building these models must thus be re-evaluated in the light of what is known about ape and fossil vocal anatomy.

A second problem with existing models of the evolution of speech anatomy concerns its relation to models of the biological and cultural evolution of communication, i.e., with external validation through model sequencing. Even if we could establish a sequence of vocal tracts, leading from ape-like to human-like shapes in gradual steps, that in itself, although an important step, would not provide an evolutionary explanation. As we and others argued elsewhere (e.g., Parker & Maynard Smith, 1990; Zuidema & de Boer, 2003, 2009), evolutionary explanations must provide a 'path of ever increasing fitness', where every new variant provides a fitness advantage in a population where the previous variant is still common. In the case of vocal tract evolution, it is unclear what the appropriate fitness function is. Existing models tend to assume that it is a simple function of the size of the acoustic space allowed by a particular vocal tract configuration. But fitness due to speech must be a function of how well an individual communicates with others in a population, which in turn depends on the communication system the population uses. However, the relation between the repertoire of speech sounds that emerges in a population and the anatomical and neurocognitive features of individuals is far from trivial.

Models that study the emergence of such repertoires have focused on vowel inventories, and on a role for self-organization in shaping them (Glotin, 1995; Berrah & Laboissière, 1999; de Boer, 2000a; Oudeyer, 2005), given constraints on the vowel space formalised by existing models of vowel perception and production. This group of models is a good example of model parallelization: different models all show the emergence of similar phenomena. They are not a good example of model sequencing, however: although these models have yielded a beautiful connection between empirical data on vowel systems and biophysical constraints, it is clear that they only scratch the surface of the full set of phonological universals: they have, for instance, little to say about consonants, syllable-structure or supra-segmental speech patterns.

Ultimately, the connection between phonology and anatomical and neurocognitive features needs to become clear to allow us to evaluate particular scenarios of the evolution of speech. However, despite the progress in modelling vocal tract evolution and vowel universals, we're still quite far from a model-based understanding of the evolution of speech. In the required sequence of explanatory models we still observe, for a variety of reasons, many gaps.

One reason is that, when addressing these more complex issues, the limits of what is at present possible with computer models are reached quickly. It is then tempting to use high-level abstractions (such as distinctive features, constraints and rule-based phonological explanations). However, making use of such abstractions, which have after all been derived for description of modern human language, and are in general not based on direct observation of neurocognitive mechanisms, incurs the risk of implicitly including the phenomena to be explained in the model – and thus resorting to phenomenological rather than explanatory modelling. For example, from typological studies it is known which consonants are unusual (for example uvular plosive [q]) and which are common (for example velar plosive [k]), but there is no language-independent biophysical and neurocognitive model that reliably predicts which articulations are more difficult to produce than others. Thus research into more complex aspects of speech is not only hampered by the computational complexity of such models, but also by our lack of knowledge about the underlying phenomena.

Likewise, we have no models of the evolution of the vocal cords. Although there are many models for human vocal cords (Dudgeon, 1970; Ishizaka & Flanagan, 1972; Titze, 1973, 1974, 2008) and some models of the interaction between the vocal cords and the vocal tract (Flanagan & Meinhart, 1964; Titze, 2002, 2008) as far as we are aware, no models exist of either chimpanzee vocal cords or of hypothetical ancestral vocal cords. This has undoubtedly to do with the lack of anatomical data (although some has recently been presented Demolin & Delvaux, 2006) but also with the fact that vocal cords (and their interaction with the vocal tract) are much more difficult to model than the acoustics of the vocal tract itself.

Another reason is that in spite of much parallel modelling effort, in some domains no consensus is reached. There is, for example strong controversy in the study of the articulatory abilities of Neanderthals and the role of modern human vocal anatomy (with its descended larynx). In this debate, Lieberman (Lieberman & Crelin, 1971) and Carré et al. (Carré *et al.*, 1995) propose that vocal anatomy has evolved for speech, while Boë et al.

(2002) propose that it has not evolved for speech, because (neural) control is more important. They reach opposite conclusions, even though they use very similar modelling techniques. The debate has lead to a rather heated exchange (Boë *et al.*, 2007; Lieberman, 2007).

Finally, some topics seem to be simply overlooked. For instance, important innovations in the cognitive adaptations for using speech that occurred between the LCA and modern humans have not been addressed by modelling. These include the ability to productively use combinatorial structure of speech and the (related) ability to learn large sets of complex utterances. Such models would be quite complex computationally, but their results might be transferable to other aspects of language, most notably syntax. After all, it has been proposed that the sequential processing and learning that are necessary for using syntax are based on adaptations for the sequential processing and learning mechanisms that are necessary for using combinatorial utterances (Carstairs-McCarthy, 1999).

Given these gaps in our understanding of the evolution of speech, the possibilities for external validation are at present limited and we should guard against over-interpreting modelling results. A case in point is the reception of Nowak *et al.* (1999), who presented an information-theoretic model and a mathematical proof of the conditions for combinatorial coding to have a fitness advantage. This proof is an elegant example of internal validation. The model fits into a larger research program in which a number of proofs of mathematical models related to the evolution of language have been presented in high-profile publications (Nowak & Krakauer, 1999; Nowak *et al.*, 2001, 2002). These models have been interpreted by other researchers as having "...demonstrated the evolvability of the most striking features of language..." (Pinker, 2000). However this confuses internal validation (the models are internally consistent) with external validation (the models correspond to reality). The latter is unfortunately far from established, given the many simplifying assumptions in Nowak et al.'s (Nowak et al., 1999) model, as we have pointed out elsewhere (Zuidema & de Boer, 2009).

4 — Conclusions

There are a number of lessons we would like to be drawn from our analysis of the state-of-the-art of language evolution modelling. First of all, it seems modellers should pay more attention to how their models relate to other models, and how they fit into particular scenarios. Although most papers on modelling the evolution of language do a good job at internal validation and at crediting other researchers' work, authors do not often make explicit which scenario they feel their model fits in and in what way their model provides external validation for other models or how other models provide it for theirs.

Second, we note that there is no lack of models and no lack of data, but there is a rather uneven distribution of modelling effort over relevant questions. It is perhaps not surprising that (as in other fields of scientific inquiry) the majority of papers are concentrated around the easiest questions. Understandable as this is, we have now reached a stage where we should also attempt to tackle the more difficult questions, and consider carefully whether a collection of models together constitute a convincing scenario.

Papers presenting 'verbal', complete scenarios can be very useful in structuring such a research program — even if we agree that one should be careful with papers that present scenarios of complex historical processes such as the evolution of language (it is all too easy to resort to speculation and wishful thinking). Jackendoff (2002) is one of the few authors who provides a rather detailed scenario that may provide a useful framework; the research field would benefit if more authors would provide a sketch of such a scenario with their work and describe how they feel their work and previous work fits in the scenario. This may help to identify the areas of language evolution that are relatively well-understood and well-studied and areas that are still *terra incognita*.

In Jackendoff's theory, a number of major transitions occurred in the evolution of language that correspond to major design features of natural language. He views, among others, an open learned vocabulary, a combinatorial phonological system, a compositional semantics, hierarchical phrase-structure and a system of syntactic categories to convey semantic relations as crucial innovations in the evolution of human language from an ancestral primate communication system. Of those five, only combinatorial phonology and compositional semantics have been addressed by multiple, parallel modelling studies.

Jackendoff's scenario is not the final word on language evolution, of course. It lacks attention for the communicative setting in which language evolved, for the possible selective advantages it offered, and to whom, and for the question why languages are so diverse and continue to change. Evolutionary biology offers many models of the evolution of altruistic traits and communication, but the relation between such models and models of the evolution of other aspects of natural language have not received much attention. Likewise, sociolinguistics and (formal) pragmatics offer many ideas about the function of language and language variation, and about the question who benefits. Also their relation to models of the evolution of language's design features remain underexplored.

From these considerations and those presented earlier in this paper, we can compose a list of key challenges for language evolution modelling that we hope will be addressed in the next few years. We present such a list in table 1; if these challenges — on the evolution of particular traits studied in the traditional subfields of linguistics and the relation to the human phenotype more broadly — are taken up by the field, we should have in a few years several models for each issue *in parallel*, as well as a set of models that *in sequence* really speak to the plausibility of a particular scenario. Only then are we approaching *external validation* of *explanatory models* of language evolution, and is the modelling approach really proving its worth to the language evolution field at large.

Table 1: Key open challenges in language evolution modelling

Phonetics & phonology:

1. Modelling the evolution of the human vocal cords;
2. Modelling the evolution of human-like (combinatorial) phonology: consonants, syllable structure, pitch/formant relation, intonation contours;

Semantics & pragmatics:

3. Modelling the transition from a closed to an open, learned repertoire of signs;
4. Modelling the evolution of duality of patterning: combinatorial phonology with compositional semantics in a unified model;
5. Modelling the evolution of human-like (compositional) semantics: quantifiers, numerals, functional/contentive split, categoricity/vagueness relation, negation;
6. Modelling dialog: how can structured, repeated communicative interactions evolve (as opposed to isolated signals);

Morphosyntax:

7. Modelling the evolution from "flat" utterances of hierarchical phrase-structure, ;
8. Modelling the evolution of word order/rich morphology trade-off;
9. Modelling the evolution of syntactic categories over and above semantic categories;

Language change & sociolinguistics:

10. Modelling the evolution of ongoing linguistic change – why are there no 'sinks' in language change?;

Relation to non-linguistic issues:

11. Language as a green beard – connection between evolution of language and altruism;
12. Language as a mental tool – connection between language and other uniquely human cognitive traits (music, consciousness, reasoning).

References

Berrah, A.-R., & Laboissière, R. (1999). Species: An evolutionary model for the emergence of phonetic structures in an artificial society of speech agents. In D. Floreano, J.-D. Nicoud & F. Mondada (Eds.), *Advances in artificial life, lecture notes in artificial intelligence* (Vol. 1674, pp. 674–678). Berlin: Springer, 1999.

Boë, L.-J., Heim, J.-L., Honda, K., & Maeda, S. (2002). The potential Neandertal vowel space was as large as that of modern humans. *Journal of Phonetics, 30*(3), 465–484.

Boë, L.-J., Heim, J.-L., Honda, K., Maeda, S., Badin, P., & Abry, C. (2007). The vocal tract of newborn humans and Neanderthals: Acoustic capabilities and consequences for the debate on the origin of language. A reply to Lieberman (2007a). *Journal of Phonetics, 35*(4), 564–581.

Carré, R., Lindblom, B., & MacNeilage, P. F. (1995). Rôle de l'acoustique dans l'évolution du conduit vocal humain. *Comptes Rendus de l'Académie des Sciences, Série II, 320*(série IIb), 471–476.

Carstairs-McCarthy, A. (1999). *The origins of complex language: An inquiry in the evolutionary beginnings of sentences, syllables, and truth*. Oxford: Oxford University Press.

de Boer, B. (2000a). Emergence of vowel systems through self-organisation. *AI Communications, 13*, 27–39.

de Boer, B. (2000b). Self organization in vowel systems. *Journal of Phonetics, 28*(4), 441–465.

de Boer, B. (2005). Evolution of speech and its acquisition. *Adaptive Behavior, 13*(4), 281–292.

de Boer, B. (2009). Why women speak better than men (and its significance for evolution). In R. Botha & C. Knight (Eds.), *The prehistory of language* (pp. 255–265). Oxford: Oxford University Press.

de Boer, B., & Fitch, W. T. (in press). Computer models of vocal tract evolution: An overview and critique. *Adaptive Behavior*.

Demolin, D., & Delvaux, V. (2006). A comparison of the articulatory parameters involved in the production of sounds of bonobos and modern humans. In A. Cangelosi, A. D. M. Smith & K. Smith (Eds.), *The evolution of language: Proceedings of the 6th international conference (evolang6)* (pp. 67–74). New Jersey: World Scientific.

Dudgeon, D. E. (1970). Two-mass model of the vocal cords. *Journal of the Acoustical Society of America, 48*(1A), 118.

Fant, G. (1960). *Acoustic theory of speech production*. 'sGravenhage: Mouton.

Flanagan, J. L., & Meinhart, D. I. S. (1964). Source-system interaction in the vocal tract. *Journal of the Acoustical Society of America, 36*(10), 2001-2002.

Gilbert, N., & Troitzsch, K. G. (2005). *Simulation for the social scientist, second edition*. Maidenhead (UK): Open University Press.

Glotin, H. (1995). *La vie artificielle d'une société de robots parlants: Émergence et changement du code phonétique*. Grenoble: DEA sciences cognitives-Institut National Polytechnique de Grenoble.

Ishizaka, K., & Flanagan, J. L. (1972). Synthesis of voiced sounds from a two-mass model of the vocal cords. *The Bell system technical journal, 51*(6), 1233-1268.

Jackendoff, R. (2002). *Foundations of language*. Oxford: Oxford University Press.

Landsbergen, F. (2009). *Cultural evolutionary modeling of patterns in language change: Excercises in evolutionary linguistics*. Utrecht: LOT.

Lieberman, P. H. (2006). Limits on tongue deformation - diana monkey formants and the impossible vocal tract shapes proposed by riede et a. (2005). *Journal of Human Evolution, 50*(2), 219-221.
Lieberman, P. H. (2007). Current views on Neanderthal speech capabilities: A reply to Boë et al. (2002). *Journal of Phonetics, 35*(4), 552–563.
Lieberman, P. H., & Crelin, E. S. (1971). On the speech of Neanderthal man. *Linguistic Inquiry, 2*, 203–222.
MacLarnon, A., & Hewitt, G. P. (1999). The evolution of human speech: The role of enhanced breathing control. *American Journal of Physical Anthropology, 109*(3), 341–343.
Nowak, M. A., Komarova, N. L., & Niyogi, P. (2001). Evolution of universal grammar. *Science, 291*(5501), 114–118.
Nowak, M. A., Komarova, N. L., & Niyogi, P. (2002). Computational and evolutionary aspects of language. *Nature, 417*(6889), 611-617.
Nowak, M. A., & Krakauer, D. (1999). The evolution of language. *Proceedings of the National Academy of Sciences, 96*, 8028–8033.
Nowak, M. A., Krakauer, D., & Dress, A. (1999). An error limit for the evolution of language. *Proceedings of the Royal Society of London, 266*, 2131–2136.
Oudeyer, P.-Y. (2005). The self-organization of speech sounds. *Journal of Theoretical Biology, 233*(3), 435–449.
Parker, G. A., & Maynard Smith, J. (1990). Optimality theory in evolutionary biology. *Nature, 348*, 27–33.
Pinker, S. (2000). Survival of the clearest. *Nature, 404*, 441–442.
Riede, T., Bronson, E., Hatzikirou, H., & Zuberbühler, K. (2005). Vocal production in a non-human primate: Morphological data and a model. *Journal of Human Evolution, 48*(1), 85-96.
Stevens, K. N., & House, A. S. (1955). Development of a quantitative description of vowel articulation. *Journal of the Acoustical Society of America, 27*(3), 484–493.
Titze, I. R. (1973). The human vocal cords: A mathematical model part i. *Phonetica, 28*(3), 129–170.
Titze, I. R. (1974). The human vocal cords: A mathematical model part ii. *Phonetica, 29*(1), 1-21.
Titze, I. R. (2002). Regulating glottal airflow in phonation: Application of the maximum power transfer theorem to a low dimensional phonation model. *Journal of the Acoustical Society of America, 111*(1 Pt 1), 367-376.
Titze, I. R. (2008). Nonlinear source–filter coupling in phonation: Theory. *Journal of the Acoustical Society of America, 123*(5), 2733-2749.
Zuidema, W., & de Boer, B. (2003). How did we get from there to here in the evolution of language? *Behavioral and Brain Sciences, 26*(6), 694–695.
Zuidema, W., & de Boer, B. (2009). The evolution of combinatorial phonology. *Journal of Phonetics, 37*(2), 125–144.

A BRIEF REVIEW OF COMPUTATIONAL MODELING IN EVOLUTIONARY LINGUISTICS

TAO GONG

Department of Linguistics, Max Planck Institute for Evolutionary Anthropology, Deutscher Platz, 6, Leipzig, 04103, Germany

Evolutionary linguistics (EL, Hauser et al., 2007) aims to identify when, where, and how a language emerges, changes, and dies out (Ke & Holland, 2006). Due to the fact that linguistic behaviors are not easily retrievable from fossil records, EL has *a multi-disciplinary nature*: knowledge, findings, and approaches from other linguistic and nonlinguistic fields can all contribute to our understanding of human language and its evolution. *Computational modeling* (CM) has been recently adopted in evolutionary linguistics. CM in EL can be viewed as the 'operational' hypotheses or theories that are expressed in computer programs and the simulation results of these programs can be viewed as the empirical predictions derived from the incorporated hypotheses or theories (Parisi & Mirolli, 2007). This abstract gives a brief review of this approach to help bridge the disciplinary gap between linguists and modelers, and the future development in CM relies greatly upon collaborations between linguists and scholars from a number of relevant disciplines.

CM is necessary in EL. CM can assist empirical studies of language evolution. Generally speaking, a model can show whether its assumptions are internally coherent, whether under these assumptions it can produce the claimed effects, and whether those assumptions are sufficient in principle to generate a given phenomena. Although a model cannot provide if those assumptions are empirically valid, it can test the available hypotheses, suggest new perspectives and help to ask more focused questions (Jäger et al., 2009). In addition, CM is an efficient way to study language as *a complex adaptive system* (Steels, 2000) and builds up a theoretical foundation for the study of language (Loreto & Steels, 2007). Human language is adaptive, consisting of multiple interacting entities that are constructed in a hierarchical way; personal experiences, social interactions, and cognitive processes could all affect language evolution. CM adopts a synthetic, bottom-up strategy, and builds up a foundation to implement scenarios that involve multiple components and their complex interactions. By isolating various components, CM can systematically analyze particular factors and their effects on language evolution, and obtain both qualitative and quantitative understanding on human language and its evolution. Such delicate control over multiple conditions is necessary for studying human language, but it is usually difficult in empirical studies. Finally, CM is validated in many aspects. Many CM studies adopt plausible assumptions supported by studies from linguistics and other fields, take objective mechanisms, and follow traceable procedures to obtain some replicable results which can be traced in linguistic data and experiments on humans (Christiansen & Kirby, 2003).

There are some problems in CM studies. Researchers may ‘over-interpret’ some simulation results or do not distinguish properly the ‘built-in’ aspects of the model and its emergent properties. CM studies may also occasionally investigate trivial matters not much relevant to the target question. Since a certain degree of simplification and specification is inevitable in CM studies, we should pay close attention not to commit these two mistakes in designing and interpreting CM studies. Meanwhile, CM studies may accidently involve *errors* (mismatches between what the developer believes a model is and what the model actually is) and *artefacts* (significant phenomena caused by accessory assumptions mistakenly considered non-significant) (Galán et al., 2009). To avoid these, we need to adopt plausible choices and refer to techniques and findings in relevant disciplines when developing a model, and conduct a systematic analysis when evaluating a model.

Many contemporary models can be classified following different criteria. For example, based on the purpose of different models (Holland, 2005), there are *data-driven models*; *existence-proof models*; and *exploratory models*. Most CM studies belong to the latter two types, which are sometimes correlated. Based on the resolution of the artificial language adopted in models (Ke, 2004), there are models in which *language is a monolithic whole* (these models usually focus on cultural transmission, language contact patterns, and other language-external factors); *subsystems are independent whole* (these models usually simulate multiple linguistic components such as semantics, syntax, phonology, and their interactions); *language is embodied in use* (these models tend to address the symbol grounding problem); and *meaning and/or form are embodied* (these models usually address the evolution of linguistic universals such as compositionality and regularity). Based on whether language users are situated in an environment and whether the communication activities use unstructured tokens or structured utterances composed of multiple tokens (Wagner et al., 2003), models at the third and fourth level of resolution can be further classified as *situated, structured models*; *situated, unstructured models*; *nonsituated, structured models*; and *nonsituated, unstructured models*. Finally, based on the adopted methods to simulate linguistic communications, models can be classified as *behavioral* (or computational) *models* that simulate the actual individual behaviors in linguistic processing and *mathematical models* that abstract these behaviors and communication into mathematical equations. Mathematical and behavioral models are mutually beneficial; the former provide the theoretical foundations for the latter that study similar phenomena, and the latter exemplify the mechanisms leading to the dynamics predicted by the former.

There are three stages in which abstract linguistic theories or hypotheses are transformed into simulation scenarios or mechanisms via many latest techniques from computer science, statistical physics, and artificial intelligence: 1) *during the setup of an artificial language, individual linguistic knowledge, and learning mechanisms*, some techniques and mechanisms, such as multi-agent system, rule-based system, artificial neural network, meaning-utterance mapping, probability matrices, strength-based competition, pattern extraction, sequential learning, and so on, can be adopted in this stage; 2) *during the implementation of a communication scenario*, many communication scenarios, such as *iterated learning* (Kirby, 1999) and various forms of *language games* (Steels, 2003), are implemented; and 3) *during the analysis of the system performance*, in which the incorporated theories or hypotheses are evaluated based on either statistical analysis and mathematical proofs, or direct or indirect comparison with the available empirical data.

Based on these stages, we can summarize the general procedures to build up a computational model to study language evolution: 1) *set up an artificial language*; 2) *define linguistic knowledge and learning mechanisms*; 3) *implement a communication scenario*; and 4) *analyze the system performance*. When evaluating a CM study, we should examine each of these aspects to obtain a clear picture of the model and a good understanding of the simulation results and the target question. When designing a CM study, we should follow these steps to clearly transform abstract theories or hypotheses into physical mechanisms or simulation scenarios, and then, systematically evaluate the incorporated theories based on the simulation results.

References

Christiansen, M. H. & Kirby, S. (2003). Language evolution: Consensus and controversies. *Trends in Cognitive Sciences*, *7(7)*, 300-307.

Galán, J. M., Izquierdo, L. R., Izquierdo, S. S., Santos, J. I., del Olmo, R., López-Paredes, A., & Edmons, B. (2009). Errors and artefacts in agent-based modeling. *Journal of Artificial Societies and Social Simulation*, *12(11)*. Available at: http://jasss.soc.surrey.ac.uk/12/1/1.html.

Holland, J. H. (2005). Language acquisition as a Complex Adaptive System. In: J. W. Minett & W. S-Y. Wang (Eds.), *Language acquisition, change and emergence: Essays in evolutionary linguistics* (pp. 411-436). Hong Kong: City University of Hong Kong Press.

Hauser, M. D., Barner, D., & O'Donnell, T. (2007). Evolutionary linguistics: A new look at an old landscape. *Language Learning and Development*, *3(2)*, 101-132.

Jäger, H., Steels, L., Baronchelli, A., Briscoe, T., Christiansen, M. H., Griffiths, T., Jäger, G., Kirby, S., Komarova, N. L., Richerson, P. J., & Triesch, J. (2009). What can mathematical, computational and robotic models tell us about the origins of syntax? In: D. Bickerton and E. Szathmary (Eds.), *Biological foundations and origins of syntax* (pp. 385-410). Cambridge, MA: MIT Press.

Ke, J-Y. (2004). *Self-organization and language evolution: System, population and individual*. Doctoral Dissertation. Hong Kong: City University of Hong Kong.

Ke, J-Y. & Holland, J. H. (2006). Language origin from an emergentist perspective. *Applied Linguistics*, *27(4)*, 691-716.

Kirby, S. (1999). *Function, selection and innateness: The emergence of language universals*. New York: Oxford University Press.

Loreto, V. & Steels, L. (2007). Social dynamics: Emergence of language. *Nature Physics*, *3*, 758-760.

Parisi, D. & Mirolli, M. (2007). The emergence of language: How to simulate it. In: C. Lyon, C. L. Nehaniv, and A. Cangelosi (Eds.), *Emergence of communication and language* (pp. 269-285). London: Springer-Verlag.

Steels, L. (2000). Language as a complex adaptive system. In: M. Schoenauer, K. Deb, G. Rudolph, X. Yao, and E. Lutton (Eds.), *Proceedings of the 6th international conference on parallel problem solving from nature* (pp. 17-26). Berlin: Springer-Verlag.

Steels, L. (2003). The evolution of communication systems by adaptive agents. In: E. Alonso, D. Kudenko, and D. Kazakov (Eds.), *Adaptive agents and multi-agent systems* (pp.125-140). Berlin: Springer-Verlag.

Wagner, K., Reggia, J. A., Uriagereka, J., & Wilkinson, G. S. (2003). Progress in the simulation of emergent communication and language. *Adaptive Behaviour*, *11(1)*, 37-69.

FLUID CONSTRUCTION GRAMMAR AND ARTIFICIAL CHEMISTRY

JOACHIM DE BEULE

IRIDIA, Université Libre de Bruxelles

Language is the product of two distinct evolutionary processes: genetic evolution and cultural evolution. It is still largely unknown what specific genetic mutations were responsible for the human capacity for language, or how these contributed to our fitness. It is known however that once they occurred, cultural evolution took over, because it works much faster (Dawkins, 1976; Kirby, 2007).

In recent years, it has become clear that the usage and learning of language and its (cultural) evolution are intricately related. This makes it difficult to make models of and understand the cultural evolution of language, because it not only requires complete models of language users, i.e. an extensive cognitive architecture including perception, conceptualization, parsing, production, learning, innovation and interaction behaviors; at the same time we also need to understand what happens when a number of such language users are put together in a population (Steels, 2000; The Five Graces Groups, 2008). This is a formidable problem. Currently, there are convincing computational models for the cultural evolution of *referential* and *lexical* languages, but not of languages that involve *complicated conceptualizations* and *grammar*.

The further advancement of this line of research will require at least the inclusion of complex, symbolic, computational conceptualization and grammar formalisms in our models of individual language users, as well as an understanding of the statistical conventionalization mechanisms that govern the usage and evolution of grammatical language in a language community. In this paper, we argue that this can be achieved by formulating and studying Fluid Construction Grammar as an Artificial Chemistry.

Fluid Construction Grammar (FCG) is a computational grammar formalism that is primarily developed for investigating the emergence of grammatical language in a community of robotic communicators. It has also become an established formalism within the fields of construction grammar and cognitive linguistics. Currently, there are FCG grammars for English, Dutch, German, French, Spanish, Russian and Hungarian. The largest FCG grammar is for English and was automatically extracted from FrameNet. It contains 10000 constructions, and can both parse and produce complicated English sentences into and from their FrameNet annotated meaning.[a] This shows that FCG is currently an ideal component for building models of the cultural evolution of grammatical language because (1) FCG is open ended, and not tailored towards any specific (type of) language: it supports both dependency and phrase structure rules, and it does not restrict the type nor the number of semantic categories etc. (2) FCG supports both the automatic parsing and the generation of sentences with the same set of grammatical constructions and (3) FCG is in line with usage based theories of language and semantic frameworks like FrameNet.

Nevertheless, FCG is only one component of a complete language user, and it still needs to be embedded in a larger cognitive architecture. Another component is the learning component, which is needed to acquire a language. It determines how the behavior of a single language user will change in response to the behavior of other language learners in the community. Only certain learning behaviors support the evolution of a shared language. Currently however, such behaviors are only understood for relatively simple, mostly lexical languages. Recently, some advances in this respect were made by looking at language as a chemical system. In this paper, we will show how FCG constructions indeed resemble complex molecules in chemistry, and how the processing of language in FCG can be described in terms of reactions between them. We will also show how this can help to unravel the conventionalization mechanisms that govern the cultural evolution of grammatical language.

References

Dawkins, Richard (1976) The Selfish Gene.

Kirby, Simon (2007) Innateness and culture in the evolution of language. PNAS.

Steels, Luc (2000) Language is a Complex Adapive System. Proceedings of Parallel Problem Solving from Nature VI, Berlin, Germany.

The Five Graces Group (2008) Language is a Complex Adaptive System. SFI Working Papers.

[a] One example sentence for which the FrameNet/FCG analysis can be found was *"The river forms a natural line between the north and south sections of the city"*.

MODELLING SOCIAL STRUCTURES AND HIERARCHIES IN LANGUAGE EVOLUTION

MARTIN BACHWERK & CARL VOGEL

Computational Linguistics Group
School of Computer Science and Statistics
Trinity College, Dublin 2, Ireland
{*bachwerm,vogel*}*@tcd.ie*

Experiments conducted with models of different social structures (varying subgroup interactions and the role of a dominant interlocutor) suggest that having a dominant, interconnected member of the group can be particularly advantageous for the emergence of language.

1. Introduction

Johann Gottfried Herder, in the 18^{th} century, won a prize for his "Treatise on the Origin of Language" addressing the question of whether human-kind could have invented language, arguing that to humans, as a social species, the invention of a language was simply natural and at the same time, essential. While much has changed in the scientific analysis of language evolution since Herder, with advances in biology, neuroscience, anthropology, among other disciplines, slowly beginning to provide an empirical basis, there remains a strong following of the notion that language must have evolved as a consequence of human sociability.

One of the most prominent theories based on social traits of early humans is the Social Brain Hypothesis, advocated by Dunbar (1997). Unfortunately, this theory has been supported by only a limited amount of computational data, owing to the complexity of building a proper model of all the social and environmental interactions involved in the evolutionary process. However, while it might be an insurmountable task at this point to model any given theory in its entirety, one can nevertheless start by producing highly abstracted, yet potentially insightful partial models that can help guide the development of the underlying theories and future modelling approaches.

2. Modelling Approach

The experiments presented in this paper have been performed using the Language Evolution Workbench (LEW, Vogel and Woods (2006)). This workbench provides over 20 adjustable parameters and makes as few *a priori* assumptions as possible.

The nevertheless assumed cognitive and social abilities of agents have been motivated by several proposed solutions to the question of language evolution and fit well with a number of models, as described below.

To begin with, the agents in the LEW are equipped with the ability to observe and individuate events, i.e. an abstracted sensory mechanism. Each agent individuates events according to its own perspective, as likely as not distinct from that of companions. In order to model the communication between agents, they are assumed to be able to join in a shared attention frame around an occurring event and engage in an interaction, whereby one of the agents is assigned the intention to comment on the event and the other, listening agent understands that the observed utterance is the speaker's comment on the event and attempts to decode the meaning of the perceived symbols accordingly. These cognitive skills of attention sharing and intentionality perception have been marked as integral to the origins of language by the theory proposed by Tomasello (2003) among others.

Three further assumptions are relevant to the symbol production and perception during interactions between agents: that agents are able to produce discernible symbols at all, that such phonemes can be combined to invent new symbols and that the transmission of symbols and phonemes occurs without noise; however, agents do not necessarily segment symbol sequences identically. While such symbols are called phonemes in the simulation, there is no reason why these should not be representative of gestural signs. The first two assumptions are made on the grounds that language could not have possibly evolved without some sort of symbols being emitted. However, the physiological considerations behind the production of symbols is not a part of this model.

The LEW fits with the so called faculty of language in the narrow sense as proposed by Hauser, Chomsky, and Fitch (2002) in that the agents are equipped with the sensory, intentional and concept-mapping skills at the start, and the simulations attempt to provide an insight into how these could be combined to produce a communication system with comparable properties to a human language. Further, the LEW agents can be seen as having completed steps 1 and 2 in the accounts presented by Jackendoff (1999) or Carstairs-McCarthy (1999), i.e. autonomously re-using and inventing new symbols from a generative unit, the phoneme.

3. Experiment Design

The goal of the presented experiments was to observe the effect of different hierarchies and social structures on the overall speed and success of communication within a group of agents. This approach extends the LEW in a way that would enable it to be used at least as a partial model for the theories of the origins of language that are based on social interactions of early humans. In particular, the experiments should provide empirical data for the possibility of language emerging in differently organized social groups, building on the comparative research by Kudo and Dunbar (2001), and the effect of being organized in a 'democratic' or an ‘oligarchial’ power structure, as proposed by Gärdenfors (1993). A special case of oligarchy—dictatorship—is approximated: there is never a *semantic arbiter*.

For the current experiments, all of the parameters of LEW have been kept fixed, except for two: the ratio of group sizes and the interaction strategies for the particular groups. The selected settings represent different social structures, with the single agent usually being referred to as ‘male’, and the agents of the larger groups, as ‘females’. In total, five different settings have been selected, with 600 runs being executed for each of these and a total of 200 rounds of 10 interactions within each such run. The observed group sizes and interaction strategies were:

1. Nine females in one group, and
   (a) all interactions among females within this group;
   (b) one male outside the nine—20% of interactions include the male;
   (c) one male outside the nine—50% of interactions include the male.
2. Nine females divided into three equal groups of three, and
   (a) 80% of interactions are within-group, 10% each with other groups;
   (b) one male outside any group—50% of female interactions are within-group, 30% with the male, 10% with each of both other groups.

4. Results

The distributions of understanding success rates for 600 runs of each condition are presented in Figure 1 and suggest that there is a strong difference in the potential of language evolving in a particular group depending on the group’s hierarchical and social structure. Setting (1a) is essentially the current baseline of the LEW as all previously conducted experiments included a single group and equal chances of any agent interacting with any other agent within this group. Settings (1b) and (1c) appear to have a slightly negative effect on the communicative success within the population; it follows that the addition of a single interconnected agent, or ‘male’ is not necessarily advantageous for the whole group. Most plausibly, this can be explained by the lack of opportunity for the agents to build a common language as they are in effect too occupied with attempting to communicate with the ‘male’. However, since the chances of the ‘male’ being selected as the speaker are equal to those of any other agent, he does not have sufficient power to actually regulate and stabilize the language of other members within the group.

The two remaining settings (2a) and (2b) suggest that having a higher number of small cliques within a group of agents has a strong positive effect on the communicative success within the population. In the setting (2a), with no ‘male’ present, one observes not only a rise in the mean of the understanding rate above 35%, but also a number of outliers nearing the 50% mark. Adding a ‘male’ to this divided ‘harem’ has once again a distinctly negative effect on the communicative

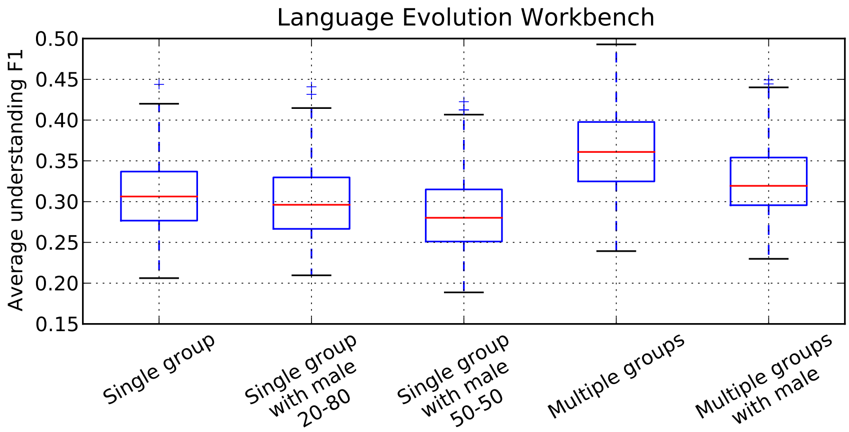

Figure 1. Distributions of understanding F1 for the five simulation settings.

success of the group. We can thus conclude with the assertion that the multiple-group scenario is evolutionary more beneficial for the emergence of language.

References

Carstairs-McCarthy, A. (1999). *The Origins of Complex Language: An Inquiry into the Evolutionary Beginnings of Sentences, Syllables, and Truth.* Oxford University Press.

Dunbar, R. I. M. (1997). *Grooming, Gossip and the Evolution of Language.* Harvard University Press.

Gärdenfors, P. (1993). The Emergence of Meaning. *Linguistics and Philosophy*, *16*(3), 285 – 309.

Hauser, M. D., Chomsky, N., & Fitch, W. T. (2002). The Faculty of Language: What Is It, Who Has It, and How Did It Evolve? *Science (New York, N.Y.)*, *298*(5598), 1569–79.

Herder, J. G. (1772). *Abhandlung über den Ursprung der Sprache.*

Jackendoff, R. (1999). Possible stages in the evolution of the language capacity. *Trends in Cognitive Sciences*, 272–279.

Kudo, H., & Dunbar, R. (2001). Neocortex size and social network size in primates. *Animal Behaviour*, *62*(4), 711–722.

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

Vogel, C., & Woods, J. (2006). A Platform for Simulating Language Evolution. In M. Bramer, F. Coenen, & A. Tuson (Eds.), *Research and development in intelligent systems* (pp. 360–373). London: Springer.

SIMULATING THE EVOLUTION OF COMBINATORIAL PHONOLOGY

W. GARRETT MITCHENER

For the workshop *Models of language evolution: Does the math add up?* at Evolang 2010.

1. Background and problem statement

The position paper by de Boer and Zuidema discusses how mathematical models and simulations are being used to develop and test theories about how the human language faculty evolved. They point out that there are significant gaps in the research community's work in this area. Primarily, there seems to be a need for modeling of biological machinery at intermediate levels of abstraction, and consequently, new simulation and modeling tools at those levels. I will discuss some of my simulations designed to fill some of these gaps. This project is in its early stages, but the preliminary results are promising.

As the position paper explains, certain parts of the overall research program in language evolution have met with great success, assisted by simulations and mathematical models. As a prime example, Nowak and his collaborators (Nowak, Plotkin, & Jansen, 2000; Plotkin & Nowak, 2000; Nowak, Krakauer, & Dress, 1999) develop an information-theoretical constraint on the lexicon and on phonology: In flat phonology, each signal requires a distinct phoneme. In combinatorial phonology, each signal is represented by several phonemes in sequence, or simultaneous phonemes transmitted over parallel channels. If there are too many topics, there may not be enough room in the space of phonemes to assign a single phoneme to each topic without incurring excessive misunderstanding. Thus, there is clear motivation for the transition from a flat phonology to an combinatorial phonology as the number of topics of conversation increases. However, details of the transition are unknown, such as what brain modifications are required to implement combinatorial signaling and what genetic mutations are required to make those modifications. This is a case where there are accurate and realistic models of the physics of speech production, and abstract models of the evolutionary forces, but I am not aware of any intermediate models that bridge the two.

It what follows, I describe a sequence of simulation of increasing complexity that bridge some of the gaps between highly abstract mathematics, physical models, and thought experiments.

2. Descriptions of simulations

Consider a problem of taking a set of Boolean inputs and producing outputs given by a fixed function of those outputs. To state the problem biologically, consider an organism that takes various actions in the presence of chemical or physical signals in its environment, thereby gaining some benefit. Let us first consider a virtual machine that represents a simplified simulation of a reaction network of genes and proteins. Complex molecules are abstracted to bit patterns. Inputs take the form of the presence or absence of specific bit patters, corresponding to signaling molecules such as the neurotransmitters and ions released by a sense organ. Outputs are given by high or low concentrations of other specific patterns. Each gene can increase or decrease the activity level of a pattern, modeling protein production, and a gene operates only when there is a sufficient amount of the protein that matches its switch pattern. The virtual machine goes through many steps in which active genes operate on pattern counts, thereby activating some genes and deactivating others, and computes a final output.

As an initial experiment, each simulated genome is considered in isolation. The genome is presented with a set of inputs, and given points for each correct output, yielding its payoff. A selection-mutation loop then steers the population toward genomes that correctly compute the desired function. To summarize the initial runs of this simulation, certain operations, such as copying inputs from one pattern to another, are easily evolved. Logical negation and compound operations require more generations and larger genomes.

Now consider the problem of transmitting information from one creature to another. To model the evolution of combinatorial phonology, suppose that the outputs of one virtual machine are transmitted across narrow channels to the inputs of another. That is, imagine that one set of inputs encodes a meaning to be transmitted, the first machine's outputs are something like motor control signals to a vocal tract, and the other set of inputs represents recognition events from a sense organ to be processed by the listening creature. We now present these creatures with the problem of copying a set of input bits across the channel. If there are few input bits, they can be transmitted directly, but if there are many, the creatures will have to develop a code and spread the message over time. The initial runs of this simulation show that this problem is distinctly harder for the artificial world to solve, however, it is able to evolve creatures that can transmit four bits across a two bit channel over time.

The next step is to first evolve creatures that must transmit few bits across a channel wide enough to accommodate them all at once, then modify the virtual world such that ever more bits must be transmitted over the same channel. This should give insights into how combinatorial phonology develops out of a fixed-size signaling system.

The overall complexity of the simulation can be increased in stages. In the first stage, the sending and receiving creatures are identical, made from of the same genome. In the second stage, creatures from genomes selected at random are paired up for the communication task and share the payoff of success. In a later stage, each creature will be allowed to overhear conversations between other creatures to have the opportunity to learn the communal code. If at all possible, that learning process should not be built into the simulation, but rather be left to evolve. In a third stage, the channel can be based on physical models of speaking and hearing, themselves parameterized by the genome. This will lead to a bridge between Nowak et al's abstract result, and anatomical phonetics.

Eventually, I would like to use a single genome to build creatures consisting of networks of these virtual machines that each receive part of the input. Then they act together to transmit it to a receiving creature, which uses its own network to interpret the message. In this form, the virtual creatures will be complete assemblies of communication devices and accompanying behavior, evolved together.

3. Why use this approach

Other researchers have studied similar simulations, notably Cangelosi and Parisi (Parisi & Cangelosi, 2002; Cangelosi & Parisi, 1998). In their mushroom world simulations, creatures are given tasks related to identifying and processing food. Those creatures evolve the ability to complete their tasks based on signals from other creatures as well as direct information about mushrooms. They evolve compositional processing in the sense of dimension reduction of semantics: Given many bits of information about a mushroom, one speaking creature can transmit a few bits to a hearer, thereby informing the hearer of the mushroom's type. That is, the creatures evolve a system for decomposing the set of possible mushrooms into a direct product of a few features. To give another example, de Boer (Boer, 2002) describes a population evolving a discrete vowel system on historical time scales, but the vowels are represented by static formants rather than signals over time, and the main results are about how phonemes spread themselves out through learning and population dynamics.

In contrast, my project approaches the evolution of compositionality from a phonological direction. My simulation is concerned with transmitting many bits given simultaneously in their entirety over a narrow channel. Instead of semantic dimension reduction, its problem of is one of evolving the serialization and deserialization system to transmit and receive a message. The messages themselves have no meaning at this stage other than that correct transmission leads to a payoff.

The traditional artificial neural networks (ANNs) used in the mushroom world do not operate with a sense of time, so my simulation must be based on a more complicated virtual machine that does. Furthermore, my virtual machine is more similar to a biological neuron, with capabilities such as pulse generation and synchrony that are absent from traditional ANNs, so it might give greater insight into what structures to look for in the human brain that might represent composition.

If the network part of the project can be made to work, it may assist biologists in identifying how biological neurons assemble themselves as a creature grows.

My virtual machines are more time-consuming to simulate than a traditional ANN, but a single workstation can compute thousands of generations of a population of hundreds in a matter of hours or days. Thus, the increased complexity does not render the simulation computationally unfeasible.

4. Simulations and the big picture

The virtual machine I described here is designed to be more like a biological cell than is typical of such simulations. It is meant to be a tool for understanding the kinds of variation relevant to language and mental computation that mutation can discover.

Without some knowledge of the biochemistry of DNA and ontogeny, the variation available to the mutation process remains hidden, and it will be difficult to put together a realistic sequence of incremental mutations that lead to the discovery of language. Incorporating effects such as gene duplication, methylation, and nucleosome binding into an artificial life simulation involves a lot of complexity and guess work, but some attempts need to be made in this area to break the impasses discussed by de Boer and Zuidema. Likewise, the scientific community has very limited knowledge of how the brain actually represents and performs computations. It seems reasonable to simulate simplified brains of increasing realism in an attempt to aid biologists in reverse-engineering those structures and their history.

References

Boer, B. de. (2002). Evolving sound systems. In A. Cangelosi & D. Parisi (Eds.), *Simulating the evolution of language.* Springer-Verlag.

Cangelosi, A., & Parisi, D. (1998). The emergence of a "language" in an evolving population of neural networks. *Connection Science*, *10*(2), 83–97.

Cangelosi, A., & Parisi, D. (Eds.). (2002). *Simulating the evolution of language.* Springer-Verlag.

Nowak, M. A., Krakauer, D. C., & Dress, A. (1999). An error limit for the evolution of language. *Proceedings of the Royal Society of London, Series B*, *266*, 2131–2136.

Nowak, M. A., Plotkin, J., & Jansen, V. A. A. (2000). Evolution of syntactic communication. *Nature*, *404*(6777), 495–498.

Parisi, D., & Cangelosi, A. (2002). A unified simulation scenario for language development, evolution, and historical change. In A. Cangelosi & D. Parisi (Eds.), *Simulating the evolution of language.* Springer-Verlag.

Plotkin, J., & Nowak, M. A. (2000). Language evolution and information theory. *Journal of Theoretical Biology*, *205*, 147–159.

PERFORMATIVE MEANING & CLAUSE-TYPE EVOLUTION

Basic clause-type distinctions are quasi-universal across languages, but little attention has been paid so far to formally modelling the evolution of these. This paper suggests a possible combination of evolutionary models that may explain the conceptual distinction between declarative and imperative clause-types. Part of this is achieved by composition of established models. But in order to account for the evolution of performative meaning of imperative clauses, extension of the standard modeling techniques seems required.

1. Declarative vs. Imperative Clause-Types in Language Evolution

Lewis (1969) already observed that the behaviorally emerging meaning of signals in ESSs of relevant signaling games is, strictly speaking, underdetermined between a declarative meaning and an imperative meaning. The meaning of a signal in an ESS could be equated with the state(s) in which it is sent and/or the action(s) which it triggers. This ambiguity is arguably adequate for primitive communication such as animal alarm calls (cf Millikan, 1995). Still, a large majority of human languages has morpho-syntactically distinct imperative forms (van der Auwera et al., 2005). These imperative forms are individuated across languages as realizing the same (quasi-)universal clause type "imperative" which is distinct from that of a "declarative" in its *stereotypical function* (Sadock & Zwicky, 1985). The question I would like to ask here is whether there is a sequence or combination of evolutionary game models to account for clause-type evolution, as distinguished by stereotypical function.

The answer obviously depends on what we take the functional distinction of clause-types to be. Lewis suggested to tell declarative and imperative meaning apart by whether it is the sender and/or the receiver who *deliberates* (cf. Huttegger, 2007, for formal implementation). However, the linguist's criterion is mainly a semantico-functional one (cf. Portner, 2004). Intuitively speaking, whereas declaratives describe states of affairs, imperatives prescribe actions.

There are two dimensions to this distinction, both of which should be dealt with in a conceptual analysis of clause-type evolution. The first is the *semantic contrast* between different denotations: states of affairs vs. actions. The challenge is to construct a game model in which signals acquire a meaning whose denotation differs relevantly in kind. Assuming imperfectly informed players, this is rather straightforward as the model sketched in section 2 shows.

More challenging is the second dimension, the *functional contrast* between description and prescription. Imperative clauses are typically associated with a particular kind of *performative*, as opposed to *descriptive*, language use. The stereotypical function of an imperative is to *give* an order or request, and not merely to describe it. This poses a challenge to mathematical modelling because we need to combine at least three crucial aspects of parallel or sequential development: (i) behaviorally emerging meaning of signals, (ii) the possibility of cooperation, (iii) the evolution of binding norms. While the first two aspects are plausibly covered by merging or sequencing existing models, especially for the latter aspect, it seems that we need to go beyond standard game models. One possibility which I suggest in section 3 is to draw on ideas from *psychological game theory* (cf. Geanakoplos et al., 1989).

2. Individuation by Semantic Denotation

We can separate emerging clause-types by semantic denotation if look we at models in which the sender may be partially uninformed about the world state, and in which the receiver may have partial information as well. Take a set $W = \{w_1, w_2\}$ of world-states and a set $A = \{a_1, a_2\}$ of receiver actions and a set M with four initially meaningless signals. Let $T = W \times A$ be the set of states of the game. The idea is that a state fixes some binary worldly parameter (think: whether it is raining or not) and the mutually best action (think: whether the receiver should go hunting or fishing). As for information allocated to players, assume that the sender always knows at least the true world state or the best action, and that the receiver has at most a piece of conditional information (such as "if it is raining, it is best to go fishing"). For more perspicuous results only, assume that the receiver also has a default action a_3. Payoffs are given by $\mathrm{U}_{S,R}(\langle w, a_{\mathtt{i}} \rangle, m, a_{\mathtt{j}}) = 1$ if $i = j$, .75 if $j = 3$, and 0 otherwise.

It can then be shown that every ESS of this game gives rise to a unique *interpretation function* $[\![\cdot]\!] : M \rightarrow \mathcal{P}(T)$ such that, roughly speaking, the interpretation of a messages $[\![m]\!] \subseteq T$ varies either only in the world component, or only in the action component. Imperative and declarative clauses can thus be identified by what is constant in different contexts of their use, and this may reasonably be taken as their semantic denotation.

3. Performative Meaning of Imperatives

In order to tell imperatives and declaratives apart functionally, we would like to devise a model in which some signals evolve to have a performative meaning. Performative meaning is the potential, given appropriate pre-conditions, to *change the facts* in relevant ways *by mere utterance* (Searle, 1969). Imperative clauses stereotypically have such performative meaning: it is by their utterance that social obligations change dynamically, and it is by an institution of social obligations that hearers may occasionally even act against their own material interests.

To model the emergence of such performative meaning is to model how a particular *normative effect* can emerge from the use of certain signals in a population.

The normative effect itself is, arguably, a dynamic change in receiver payoffs. (For a basic explanatory model, it is not essential whether these payoff changes come about by a complicated secondary mechanism involving status, fear of retaliation, or similar.) But then, to have performative meaning emerge, we would need receiver payoffs to change in some manner under evolutionary dynamics. The problem is that in standard models what adapts to evolutionary pressures is the *behavior* of players, while their payoffs remain fixed. What is needed, therefore, is a certain *self-reflexive property* of a model to make the receiver's payoffs depend on the meaning of a signal as it emerges from aggregate population behavior.

This self-reflexive property can be implemented by adopting ideas from *psychological game theory* (cf. Geanakoplos et al., 1989), where agents' payoffs are analyzed as composed of the traditional *material payoffs* and *psychological payoffs* which may vary according to the beliefs and intentions of players. These psychological payoffs may depend on the gradually acquired meaning of signals, for instance, by following Sugden (2000) who argues that mere expectations can be normative, i.e., that we are biased towards compliance with the expectations of others. If we assume that receivers have psychological payoffs susceptible to sender expectations, and if we also assume that sender expectations are a function of current population behavior, then this yields a possible mechanism how emerging behavioral patterns can impact the receiver's payoffs for performative meaning to emerge.

Suppose we have two receiver actions $A = \{a_{\mathtt{A}}, a_{\mathtt{B}}\}$. Take two sender types and three receiver types, depending on which action is preferred: senders always want exactly one action, receivers may be indifferent. With this, distinguish six possible states: $T = \{t_{\mathtt{AA}}, t_{\mathtt{AB}}, t_{\mathtt{BA}}, t_{\mathtt{BB}}, t_{\mathtt{A?}}, t_{\mathtt{B?}}\}$, where the first index gives the sender type, the second the receiver type. Each player knows her own type but not the opponent's type. Material payoffs are defined as:

$$\mathrm{U}_S(t_{\mathtt{XY}}, m, a_{\mathtt{Z}}) = \begin{cases} 1 & \text{if } X = Z \\ 0 & \text{otherwise} \end{cases} \quad \mathrm{U}_R(t_{\mathtt{XY}}, m, a_{\mathtt{Z}}) = \begin{cases} \alpha & \text{if } Y = Z \text{ or } Y = ? \\ 0 & \text{otherwise} \end{cases}$$

Here, $0 < \alpha \leq 1$ is a parameter of the model that measures, inversely, how costly performance of an undesirable action is. (If cooperativity prevails in a population, as made plausible by a multiplicity of mechanisms (cf. Nowak, 2006), this parameter will be small.) Finally, the receiver's overall payoffs are a linear combination of material and psychological payoffs $F(\cdot)$:

$$\mathrm{U}^*_R(t_{\mathtt{XY}}, m, a, \vec{r}) = (1 - \beta) \times \mathrm{U}_R(t_{\mathtt{XY}}, m, a) + \beta \times F(\cdot)$$

For psychological payoffs to implement the sender's expectation of receiver behavior, define function $F(m, a, \vec{r})$ as the proportion of a responses to m in the current receiver population $\vec{r}$.

It can then be shown that for parameter values $\beta > \frac{\alpha}{\alpha+1}$ there are ESSs in which signals are used imperatively to change the receiver's payoffs and induce the sender-optimal action even if that is not materially receiver-optimal. Furthermore, under somewhat stronger parameter conditions, $\beta > \frac{\alpha}{\alpha + 1/3}$, the system will evolve towards one of these from a neutral state in which receivers always play their own materially preferred action. This simple model demonstrates how performative meaning could arise: initially signals do not (significantly) impact receiver payoffs, but they can evolve to do so.

Taken together, I suggest that the emergence of clause-types from functionally ambiguous signaling can be plausibly explained by a combination of standard and non-standard evolutionary models. The particular combination sketched here also suggests a concrete evolutionary path: context-driven semantic distinctions can emerge whenever agents acquire sufficient reasoning capabilities to accommodate higher order uncertainty; in parallel or in sequence, performative meaning can emerge against a background of by-and-large cooperative social practices whenever, for instance, hearer's are interested to conform to speaker expectations about population behavior.

References

van der Auwera, J., Lejeune, L., Pappuswamy, U., & Goussev, V. (2005). The morphological imperative. In M. Haspelmath, M. S. Dryer, D. Gil, & B. Comrie (Eds.), *The world atlas of language structures* (pp. 286–289). Oxford University Press.

Geanakoplos, J., Pearce, D., & Stacchetti, E. (1989). Psychological games and sequential rationality. *Games and Economic Behavior*, *1*, 60–79.

Huttegger, S. M. (2007). Evolutionary explanations of indicatives and imperatives. *Erkenntnis*, *66*, 409–436.

Lewis, D. (1969). *Convention. a philosophical study.* Harvard University Press.

Millikan, R. G. (1995). Pushmi-pullyu representations. *Philosophical Perspectives*, *9*, 185–200.

Nowak, M. A. (2006). Five rules for the evolution of cooperation. *Science*, *314*, 1560–1563.

Portner, P. (2004). The semantics of imperatives within a theory of clause types. In K. Watanabe & R. Young (Eds.), *Proceedings of SALT 14.* Ithaca, NY: CLC Publications.

Sadock, J. M., & Zwicky, A. M. (1985). Speech act distinctions in syntax. In T. Shopen (Ed.), *Language typology and syntactic description: Clause structure* (Vol. 1, pp. 155–196). Cambridge University Press.

Searle, J. (1969). *Speech acts: An essay in the philosophy of language.* Cambridge University Press.

Sugden, R. (2000). The motivating power of expectations. In J. Nida-Rümelin & W. Spohn (Eds.), *Rationality, rules and structure* (pp. 103–129). Dordrecht: Kluwer.

BAYESIAN MODELS OF LANGUAGE EVOLUTION: THE NEW UNIFYING PARADIGM?

DAN DEDIU

Max Planck Institute for Psycholinguistics, Wundtlaan 1,
Nijmegen, 6525 XD, The Netherlands
Dan.Dediu@mpi.nl

Bayesian models seem to hold many promises for understanding language evolution and change, but it is not yet clear if and how they can deliver on these. They have several notable advantages, most important being a standardized and well-understood model of acquisition/learning, but also several issues and problematic assumptions, which I will briefly discuss.

1. Introduction

A quick review of the literature dedicated to modelling language change and evolution reveals a bewildering diversity of approaches, some justified by the characteristics of the problem to be addressed, some due to historical accidents boiling down to their authors' knowledge and preferences, while for some others no such reasons are readily available.

It can be argued that "one-size-fits-all" approach to modelling such a complex field is doomed to fail and thus, particular, "locally adequate" models addressing only the relevant aspects of the problem at hand are more appropriate. But it can as well be argued that such local models are *ad-hoc*, unprincipled, and thus hard to justify and almost always under-investigated.

A very promising unifying proposal seems to be represented by the "Bayesian paradigm", where a powerful, standard and well-understood machinery is adapted to the issues facing language evolution and change.

2. Bayesian Language Agents

Language learning/acquisition can be formulated in the very general framework of Bayesian statistics (Press, 2003; Oaksford & Chater, 2007), as follows (Griffiths & Kalish, 2007; Kirby, Dowman, & Griffiths, 2007; Smith & Kirby, 2008).

Let $\mathcal{H}$ be the *universe of all possible languages* (or *hypotheses*), $h \in \mathcal{H}$, defined as appropriate for the phenomena under investigation (to simplify, we will take it as discrete). A Bayesian agent, A, entertains at any moment a *distribution* over this universe which associates a "subjective" probability (Press, 2003) to every possible language, $0 \leq p(h) \leq 1, \forall h \in \mathcal{H}$, and $\sum_{h \in \mathcal{H}} p(h) = 1$. The agent uses such a distribution to either *comprehend* or *produce* language (involving others or even itself), using model-specific conventions.

Language learning/acquisition concerns the *process of constructing* a distribution over $\boldsymbol{\mathcal{H}}$ which allows the learner agent A to communicatively function in the communicative context under consideration. A is exposed to *the data*, d, – usually a set of utterances – which embody the learner's experience with the language(s)[a]. A then uses d to *update* its distribution across hypotheses, following *Bayes' theorem* (Griffiths & Kalish, 2007; Press, 2003):

$$p(h|d) = \frac{p_{obs}(d|h)p(h)}{\sum_{h' \in \boldsymbol{\mathcal{H}}} p_{obs}(d|h')}, \forall h \in \boldsymbol{\mathcal{H}} \quad (1)$$

where $p(h|d)$ is the *posterior* (updated) probability that hypothesis h holds after exposure to d, $p_{obs}(d|h)$ is the *likelihood* of observing d were h true, and the denominator is required for normalization. $p(h)$ is the *prior* probability of h before A "saw" d and can represent the outcome (posterior probability) of a previous learning round or the result of "innate" predispositions. However, there are a number of assumptions and issues behind this approach.

The "initial" prior, $p_0(h)$, before A sees any data, is usually taken to reflect *innate or genetic biases*, produced by developmental processes, but there is no necessary implication of language-specific genes (Griffiths & Kalish, 2007; Kirby et al., 2007). However, there is no clear-cut difference between "genetic" and "environmental", "development" is not a discrete, encapsulated and teleological phase in the life-cycle of an organism and "genes" are essential to all processes at all times (Minelli, 2003; West-Eberhard, 2003; Odling-Smee, Laland, & Feldman, 2003), making such a neat dichotomy of *before* – given by the genes in the form of *biases* – versus *after* – built on this basis by the environment/culture – untenable.

This identification of *genetic biases* with $p_0(h)$ is also misleading as this is not the only parameter – and maybe not even the most important – which can conceivably be affected by genes in such a Bayesian learner (e.g., the *learning algorithm* or the likelihood function). For example, it is clear that the choice between *sampler* and *maximizer* learning algorithms makes a crucial difference to the resulting dynamics of single-agent chain *Iterated Learning Model*s (Griffiths & Kalish, 2007; Kirby et al., 2007), but this difference diminishes in more complex (and realistic) scenarios (Dediu, 2009). This suggests that generalizations from such cases are unwarranted and that the apparently obvious issue of "converting" the posterior into actual utterances-producing machinery is far from trivial.

This model's *acquisitionist* assumption[b] is instantiated by the (potential) differences between the distribution used (whatever the mechanism) by an agent to

[a] $\boldsymbol{\mathcal{H}}$ and d can – and should – refer to more than strictly linguistic information, including, for example, aspects of the social and communicative context (Enfield & Levinson, 2006).

[b] The acquirer is the locus of language change, by reinterpreting the linguistic data, leading to differences between the target and learned languages (Kirby & Hurford, 2002).

produce data, $p(.) \rightsquigarrow d$, and the posterior distribution derived by the learner from these (and, potentially, other agents') data, $p'(.|d)$. However, there is a wealth of data suggesting that the acquirer is not the only (and probably not even the most important) locus of language change (Croft, 2000; Enfield & Levinson, 2006; Ostler, 2005) and a way of addressing this is by replacing the once-for-all "learning algorithm" with a dynamic selection of hypotheses from the posterior distribution function of the communicative context and history.

Recently, Ferdinand and Zuidema (2009) have shown that if the prior is taken as the learning biases, then the learner must be *omniscient* about the possible sources of the data. They also touch on the general problem of what exactly is Bayes' theorem *supposed to represent* (Oaksford & Chater, 2007; Press, 2003): is it a description of the actual computations going on in the learner's head or is it an idealized model which is approximated to various degrees by other mechanisms and, if so, why, how, how well and what about the deviations?

The nature of the likelihood function also needs to be specified, as it seems necessary that it somehow captures the actual production mechanisms in order to result in meaningful posterior distributions. If so, how does a learner arrive at this match?

There is also the considerable computational burden of updating the probability distributions, especially when realistically large hypotheses spaces are considered, and we must face this head-on when running our computational models[c]. Conjugate priors (Press, 2003) might help us, but they are limited to very simple models and will inject a whole new set of assumptions and constraints, making it harder to understand the causes of the results.

3. Conclusions and some suggestions

As de Boer and Zuidema (2009) outline in their manifesto, there seems to be a deep dissatisfaction at several levels, and especially in what concerns the meaning of, and articulation between, various models. This could be addressed by building a public database[d] containing not only bibliographic information of relevant publications[e], but, most importantly, a fully searchable database of *models* with detailed description, optional software/source code and typical results.

I have argued above, in line with other researchers, that the Bayesian approach is probably not a general solution to our problems. But it can probably be used "as appropriate" to model local processes where strong substantive arguments can be made for the fit of a Bayesian description

To conclude, I think that, in order for the field to mature, we need an effort

[c]A Hierarchical Bayesian simple ILM chain inspired by (Dediu, 2009) using a Gibbs sampler (JAGS, http://www-fis.iarc.fr/ ˜martyn/software/jags/) took on the order of days on a top-end machine.

[d]Like the *National Center for Biotechnology Information*'s portal, http://www.ncbi.nlm.nih.gov/.

[e]As Jun Wang's essential *Language Evolution and Computation Bibliography*.

of synthesis of various competing models and the issuing of (soft) "recommendations" concerning their properties, assumptions, levels of fit and inter-connections. Of course, such an effort might seem less glamorous than creating new models and will require the collaboration of the best modelling minds, but it holds the promise that one day we would have a more coherent modelling landscape, the way (apparently) population/evolutionary genetics has.

Acknowledgements

I wish to thank K. Smith, J. Zuidema, S. Kirby, N. Enfield and A. Dima.

References

Croft, W. (2000). *Explaining language change: An evolutionary approach.* Pearson Eduction Limited.

de Boer, B., & Zuidema, W. (2009). Models of language evolution: Does the math add up? In *ILLC Preprint Series (pp-2009-49).* University of Amsterdam.

Dediu, D. (2009). Genetic biasing through cultural transmission: do simple Bayesian models of language evolution generalize? *JTB*, *259*, 552-561.

Enfield, N. J., & Levinson, S. C. (Eds.). (2006). *Roots of human sociality: Culture, cognition and interaction.* Berg Publishers.

Ferdinand, V., & Zuidema, W. (2009). Thomas' theorem meets Bayes' rule: a model of the iterated learning of language. In *COGSCI* (p. 1786-1791).

Griffiths, T., & Kalish, M. (2007). Language evolution by iterated learning with bayesian agents. *Cognitive Science*, *31*, 441–480.

Kirby, S., Dowman, M., & Griffiths, T. (2007). Innateness and culture in the evolution of language. *PNAS*, *104*, 5241–5245.

Kirby, S., & Hurford, J. (2002). The emergence of linguistic structure: An overview of the iterated learning model. In A. Cangelosi & D. Parisi (Eds.), *Simulating the evolution of language* (p. 121-148). Springer Verlag.

Minelli, A. (2003). *The development of animal form: Ontogeny, morphology, and evolution.* Cambridge University Press.

Oaksford, M., & Chater, N. (2007). *Bayesian rationality: The probabilistic approach to human reasoning.* Oxford University Press.

Odling-Smee, J. F., Laland, K. N., & Feldman, M. W. (2003). *Niche construction: The neglected process in evolution.* Princeton University Press.

Ostler, N. (2005). *Empires of the word: a language history of the world.* Harper Collins Publishers.

Press, S. (2003). *Subjective and objective Bayesian statics.* John Wiley & Sons.

Smith, K., & Kirby, S. (2008). Cultural evolution: implications for understanding the human language faculty and its evolution. *Phil Trans R Soc B*, *363*, 3591–3603.

West-Eberhard, M. J. (2003). *Developmental plasticity and evolution.* Oxford University Press.

COMPUTATIONAL MODELING OF LANGUAGE EVOLUTION: GAPS AND CHALLENGES

MONOJIT CHOUDHURY
Microsoft Research Lab India
Bangalore, 560080, India

ANIMESH MUKHERJEE
Complex Systems Lagrange Lab, ISI Foundation
Torino, 10133, Italy

1. Computational Models of Language in Tinbergen's Framework

Language evolution is arguably one of the hardest problems in science and is highly interdisciplinary in nature (Christiansen and Kirby, 2003). As evident from the exponential growth in number of publications, this area of research, like other scientific disciplines, has popularly adopted the use of mathematics and, more recently, computational techniques. Nevertheless, modeling seems to be quite a hard problem in language evolution and is still in its infancy (de Boer and Zuidema, 2009). In order to understand the difficulties faced by this emerging sub-discipline we advocate that one should not only look into the challenges and issues within language evolution, but rather *language* itself as the topic of research. In fact, by doing so, as we shall see, one can arrive at two important conclusions – (a) knowledge sharing between different communities studying language is currently limited even though cross-fertilization of ideas seems extremely necessary for the progress of the field, and, (b) there is a need for data creation and analysis before one jumps into modeling.

Language is a human behavior and therefore, it is reasonable to organize language research within the framework of *Tinbergen's four questions* (Tinbergen, 1963). In Table 1, we lay out the areas of language research within these four questions, and then list out the scope of mathematical and computational modeling within each division. We also note the current availability of data for each kind of research and ease of gathering more data.

Table 1. Understanding the interrelations between various disciplines that study "language" and its computational models within the framework of Tinbergen's four questions for ethology.

Tinbergen's questions	*Research Areas*	*Computational Models/Methods*	*Availability of data*
Causation	a) Neuropsychological processing of language b) Physiology of speech production and perception	a) Cognition aware models of NLP (esp. parsing) b) Speech technology (esp. parametric models)	Moderate to high; Easy to gather
Ontogeny	a) Language acquisition b) Development of other communicative traits	a) Machine learning b) Models of human learning	Little to moderate; Moderately easy to gather
Adaptation	a) Uniqueness of human language b) Language universals and typological studies c) Relevant biological adaptations	a) Quantitative and statistical properties of languages b) Findings from NLP might be of use	High; Quite easy to gather
Phylogeny	a) How have languages evolved b) How do languages change? c) Evolution of brain and speech apparatus	a) Models of language evolution & change: math and simulation based b) Cladistics and computational phylogenetics	None to little; Very hard or almost impossible to gather

Computational Linguistics (CL) or Natural Language Processing (NLP) also study and develop computational models of language. However, the focus is on human language technology, e.g., machine translation and speech processing (Jurafsky and Martin, 2009), and therefore, they do not fit into Tinbergen's framework. On the other hand, some of these research areas can really benefit and get benefitted by the research in other areas of language (e.g., speech technology and physiology of speech production and perception, machine learning for NLP and language acquisition). We do not see, however, much exchange of ideas between the language evolution or cognitive science communities and the NLP community.

In general, a careful study of Table 1 reveals two important points – (a) there is very little sharing of data and knowledge between the various communities studying computational models of language; and (b) there is very little data for directly modeling the problems within "phylogeny". However, other areas have access to large amount of data or at least possibility of gathering data. We argue that these factors, along with the fact that there are not many known or accepted ground truths to base the successful models of language evolution constitute the major hurdles in the area of language evolution. In the following three sections we discuss these issues briefly.

2. Lack of Knowledge Sharing

While there is a very vibrant and huge research community working on NLP, there is very little sharing of knowledge between them and the language evolution community. This is apparent from the facts that (a) there are hardly any cross-citations between these communities, and (b) we hardly see works on language evolution being published in ACL conferences. Fragmentation of the community has not only ceased cross-fertilization of ideas, but also made the communities very small. Any research area needs a *critical mass* of researchers to prosper (Shneider, 2009). We believe that lack of critical mass is one of the most detrimental factors in the area of language evolution. This also explains some of the other issues raised in (de Boer and Zuidema, 2009), such as why there has been a lot of modeling efforts only in few areas, even though there is data in the other areas. Indeed, there are not enough researchers in the community to work on many different data sets.

As an aside, it is interesting to note that in many other scientific disciplines, for example physiology and medicine, there is a strong sharing of knowledge. Physiological findings go into development of drugs and surgical procedures, whereas clinical data feed into the models of physiology.

3. Lack of Data

There is a lot of linguistic (e.g., corpora, treebanks, phonological databases) and psycholinguistic data (CHILDES, semantic relatedness of words) available for research. In certain areas, where data is not available, it might not be hard to gather more data using advanced techniques and sophisticated instruments. Nevertheless, we have very little access to data for directly validating models of language evolution and it is highly unlikely that we will ever be able to gather large quantities of such data. Therefore, one has to be content with indirect validation. Since insufficient data cannot distinguish between the good and the bad models, limited access to data prevents direct validation of the models leading to a loss of credibility.

4. Lack of Ground Truths and Accepted Frameworks

There is hardly any consensus on theories of language evolution; neither do we have concrete theories of language acquisition and processing. In fact, the most celebrated “Chomskyian” framework of *principles and parameters* is also contested. Therefore, it is hard to come up with a set of constraints or a framework which a model should follow. In absence of any such principles, the space of possible models become infinite and lack of data for validation makes it harder to make strong claims based on computational modeling.

Note that, let alone physical sciences, even biological sciences have well accepted frameworks (Darwinian Theory), ground truths (central dogma of molecular biology) and quantitative laws (Junck, 1997).

5. Conclusions

Every scientific discipline goes through several stages of evolution (Shneider, 2009), and we believe that language evolution and especially the modeling approach is in the *first stage*, which is the time for spelling out the agenda, defining the vocabulary and coming up with the basic principles. Therefore, this area of research will benefit from some of the suggestions, such as identification of relationship between models and gaps therein, put forward by de Boer and Zuidema (2009). Nevertheless, we do feel that there is an urgent need for – (a) knowledge sharing with related communities and especially the NLP community which already has gathered lot of data and useful techniques; this can be achieved through organization of conferences and popularizing the work in various venues, (b) concentrating on creation and analysis of data, rather than jumping into synthetic and explicatory models. There are numerous examples of synthetic models in this area, which fall flat if the data is investigated a little deeper, and, (c) laying out the principles of modeling, such as the constraints that any model should satisfy, and the minimum needs for validation. The community should be encouraged to report the failed models as well.

References

Christiansen, M. H., and Kirby, S. (2003) Language Evolution: The Hardest Problem in Science? In *Language Evolution: The States of the Art*. OUP.

de Boer, B., & Zuidema, W. (2009). Models of Language Evolution: Does the Math Add Up? *ILLC Preprint Series* (PP-2009-49)

Junck, J. R. (1997). Ten Equations that Changed Biology: Mathematics in Problem-Solving Biology Curricula, *Bioscene*, **23**, pp. 1-36

Jurafsky, D., and Martin, J. H. (2009). *Speech and Language Processing: An Introduction to Natural Language Processing, Speech Recognition, and Computational Linguistics*. 2nd edition. Prentice-Hall.

Shneider, A. M. (2009). Four Stages of a Scientific Discipline; Four Types of Scientist. *Trends in Biochemical Science* **34**(5), pp. 217-23

Tinbergen, N. (1963). On Aims and Methods in Ethology. *Zeitschrift für Tierpsychologie*, **20**, pp. 410 – 433

MODELING LANGUAGE EVOLUTION: A GENERAL FRAMEWORK WITH THE MIRROR SYSTEM HYPOTHESIS AS A SPECIAL CASE

MICHAEL A ARBIB

University of Southern California, Los Angeles, USA

While stressing the importance of calibrating focused models of specific phenomena against each other and available data, De Boer & Zuderna (2009) also note the value of 'verbal', complete scenarios for the evolution of language, citing Jackendoff (2002) as providing one of the few such scenarios. De Boer & Zuderna concludes with a Table listing key open challenges in language evolution modeling. However, none of the listed challenges are explicitly related to brain mechanisms. In response to this, I offer an alternative general framework which makes explicit a range of choices which set the general nature of a model even before a range of data is chosen for consideration (or, conversely, which reflects a strong restriction on the range of data to be considered relevant). The paper then presents the specific choices made in a brain-centered research program which has evolved [sic] from the Mirror System Hypothesis, seeking to calibrate a scenario for the evolution of language against new approaches to neurolinguistics.

1. A Diversity of Divides in Devising System Models for Language Evolution

Chomsky (1956) showed that certain classes of grammars align with certain classes of automata – most famously, the alignment between context-free grammars and push-down automata. Over the years, Chomsky has moved on from this highly formal framework to seek to characterize a class of grammars broad enough to encompass the syntax of all human languages, with the resultant framework changing quite drastically every decade or so. Chomsky’s writings introduce another automaton, the Language Acquisition Device (LAD), but this time without a precise formal framework. It is posited to be innately specified, and present in (almost) every human infant.. Its task of is to accept strings from a single language L as input and to infer a grammar G(L) which generates all and only the strings of L. Most of Chomsky’s writing discounts the notion that there can be a useful account of the evolution of language (see Hauser et al. 2002, for an apparent exception -- but one that seems inconsistent with much that has gone before). But for those who wish to study language evolution in the Chomskian framework (e.g., Pinker & Bloom 1990), the study of evolution reduces to breaking Universal Grammar into several components that can be placed within an evolutionary succession. On this view, the available data are simply well-formedness judgments for each language under study. However, working backwards in time, we can see that different approaches to language evolution may consider language in a broader sense of performance that includes:

- Data on language use in communication and thought, including use in conversation, including a whole range of speech acts.
- Data on language acquisition.
- Data from historical linguistics.
- Archeological data which seek to infer clues to the language of a culture from the material traces of that culture.
- Animal communication systems in general, seeking commonalities with and differences from language.
- Animal behavior in general, seeking commonalities with and differences from communication systems.
- Data on brain function.

Our general concern, then, is: how did evolution equip *Homo sapiens* to enable human societies to develop a range of languages; and an individual to acquire the language of the community in which it is raised? The evolutionary challenge is both biological and cultural and there are many divisions over what the nature of this challenge might be:

Biological: How did biological evolution endow* Homo sapiens *with brains and bodies that can acquire and use language?

- *Modeling Divide* 1: Our ancestors evolved a capability for protolanguage – which had an open lexicon but little if any syntax – before they developed language *versus* no intermediary was involved.
- *Modeling Divide* 2: Models must be speech-centered, equating speech with language, *versus* we seek to understand the shared mechanisms that support signed languages as well as spoken languages.
- In any case, the importance of spoken language asks us to understand the evolution of the vocal apparatus which made it able to produce speech, and the related mechanisms for perception, and for the neural control of perception and production.
- *Modeling Divide* 3: Biological evolution of language must make crucial use of data on the brain *versus* the primary data are those of linguistics alone.

Cultural: What aspects of language are innate, and what are the fruits of historical change?

- *Modeling Divide* 4: Biological evolution endowed us with a Universal Grammar *versus* it gave us mechanisms which made the eventual *invention* of language possible.

- In either case, one needs to understand what the child learns during language acquisition, and what biological evolution provided to make such learning possible.
- *In either case,* one needs to understand what happens during historical language change, and what biological evolution provided to make such processes possible.

Modeling Divide 5: Initially, much of protolanguage was holophrastic (with a protoword describing a frequently occurring or significant situation) *versus* protolanguage started with words akin in scope to modern words (such as nouns and verbs)

Modeling Divide 6: Language evolution is to be understood solely in terms of adaptive pressures for communication or thought *versus* language evolution rests in part on the exaptation of adaptations that are not directly related to communication.

These six modeling divides (and, of course, there are others) define 64 overall approaches to language evolution. Thus any *general* framework must justify (at least) which side of each of the six divides it lies on. By contrast, more focused models may ignore many of these issues to address the possible evolution of mechanisms responsible for some specific set of data, such as how the human speech apparatus can produce the observed sounds of human languages, though even here assumptions from some general framework may play a crucial role.

Returning to our general discussion, then, our choice of behavioral data determines whether we look at *Performance Systems, Developmental Systems, Historical Systems,* and *Evolutionary Systems*. The full talk will outline the different modeling challenges these pose, and how each relates to the later types.

2. The choices made in developing the Mirror System Hypothesis

The Mirror System Hypothesis (Arbib 2005, Arbib & Rizzolatti 1997, Rizzolatti & Arbib 1998) offers an approach to language evolution rooted in the data of primatology, comparative neurobiology, and neurolinguistics. In terms of the six Modeling Divides (MDs) presented in the previous section, the Mirror System Hypothesis makes the following choices:

- MD1: Our ancestors evolved a capability for protolanguage before they developed language.
- MD2: Models must seek to understand the shared mechanisms that support signed languages as well as spoken languages.
- The importance of spoken language asks us to understand the evolution of the vocal apparatus which made it able to produce speech, and the related mechanisms for perception, and for the neural control of perception and production.

- MD3: Biological evolution of language must make crucial use of data on the brain.
- MD4: Biological evolution gave us mechanisms which made the eventual invention of language possible.
- One needs to understand what the child learns during language acquisition, and what biological evolution provided to make such learning possible.
- One needs to understand what happens during historical language change, and what biological evolution provided to make it possible.
- MD5: Initially, much of protolanguage was holophrastic.
- MD6: Language evolution rests in part on the exaptation of adaptations that are not directly related to communication.
- What is the social structure that brings the child into the community using a specific language, and what neural capabilities are needed in the brains of the child and its caregivers to support this process?

The full paper will summarize where the Mirror System Hypothesis has offered a computational model or a conceptual model for a subtopic within the framework, outlining the modeling approach used and what data have been addressed in each case. It will also summarize where models from other groups serve to fill in gaps in our own research within this general framework. It will then suggest the strategies this implies for comparing the various models with those that fall on the other side of one or more divides.

Acknowledgement: The present research was supported in part by the National Science Foundation Program in Perception, Action, and Cognition under Grant BCS-0924674 (Michael A. Arbib, Principal Investigator).

References

Arbib, M. A. (2005) From Monkey-like Action Recognition to Human Language: An Evolutionary Framework for Neurolinguistics (with commentaries and author's response). *Behavioral and Brain Sciences 28*: 105-167.

Arbib, M. A. & Rizzolatti, G. (1997) Neural expectations: a possible evolutionary path from manual skills to language. *Communication and Cognition 29*: 393-424.

Chomsky, N. (1956) Three models for the description of language. *IEEE Transactions on Information Theory 2*(3): 113--124.

de Boer, B. & Zuidema, W. (2009) Models of Language Evolution: Does the Math Add Up? *ILLC Preprint Series, University of Amsterdam (PP-2009-33)*: http://staff.science.uva.nl/~jzuidema/mole2010/deboerZuidema2010mole-workshop.pdf.

Hauser, M. D., Chomsky, N. & Fitch, W. T. (2002) The Language Faculty: What is it, who has it, and how did it evolve? *Science 298*: 1568-1579.

Jackendoff , R. (2002) *Foundations of Language: Brain, Meaning, Grammar, Evolution*, Oxford and New York: Oxford University Press.

Pinker, S. & Bloom, P. (1990) Natural Language and Natural Selection. *Behavioral and Brain Sciences 13*: 707-784.

Rizzolatti, G. & Arbib, M. A. (1998) Language within our grasp. *Trends in Neurosciences 21*: 188-194.

COMPARING MODELS OF LANGUAGE EVOLUTION WITH THE BRAIN OPERATION DATABASE (BODB)

MICHAEL A ARBIB

University of Southern California, Los Angeles, USA

Models related to language evolution may address very different topics (e.g., speech perception versus neural localization of syntax). It thus becomes infeasible to provide an a priori set of benchmarks for modeling. As a result, many modelers will simply choose a limited set of data (which may be imagined rather than empirical) and offer a model which conforms, more or less, with those data, with little if any assessment of how the model compares to other models. A similar situation (though perhaps less severe) holds in systems and cognitive neuroscience. The present paper introduces the Brain Operation DataBase (BODB, pronounced “Beau D-B”) which was designed to facilitate the calibration of focused neuroscience models of specific phenomena against each other and available data. The full talk will suggest how the tools it offers may be extended to serve models of language evolution, and then will invites readers to join in using BODB to document and compare models in the language evolution domain.

1. Introducing the Brain Operation Database (BODB)

We explain how BODB supports the documentation *and comparison* of neuroscience models of specific phenomena against each other and available data. The key point (see figure below) is that BODB supports *automatic benchmarking* – once a set of models is chosen, BODB automatically generates a set of benchmarks against which the models can be compared. This is important in a field as heterogeneous as neuroscience (or language evolution) in which no global set of benchmarks applies across the diversity of the models.

BODB provides a Model Entry Form which enables one to enter all details pertinent to a model which is being documented. One can then search the database for relevant entries that need not exist but can simply be linked. Each model must be given a **Title** as well as a **Brief description**. The **Architecture** is given first by a **Diagram** (or **Diagrams**): Since the model may be hierarchical, one may choose to use several diagrams with at least one to present the overall structure of the Model. In some cases, a diagram will explicitly show not only the immediate submodules but also the further decomposition of some of these.

A **Tour of diagram** briefly mentions all the modules and makes explicit how the interactions between them contribute to the function of the model.

The **Narrative** reviews how the model works and what it does, in relation to Summaries of Experimental Data (SEDs) and Simulation Results (SSRs) presented *and their relationships*.

Summaries of Experimental Data (SEDs) – in the neuroscience domain – summarize data gathered by neurophysiologists, psychologists, anatomists, brain imagers, etc., about the actual behavior or brain function and structure of real animals and humans. An SED should be detailed enough to explain aspects of the design of the model (Support) or test the model (Explanation or Contradiction) by explicit comparison with **Summaries of Simulation Results (SSRs)** which (no surprise) summarize the results obtained by running the model with one or more protocols. (In some cases, the SSR will summarize results of mathematical investigation of the model, rather than simulations.) A BODB entry must make explicit the relation of the SED to the model as follows:

scene setting: Some data are not actually used either in defining or testing the model but may strengthen the Model Entry by "setting the scene."

support: SEDs that *support* some aspect of model design. Documenting the relevance forces one to sharpen the presentation of the SED since it has to contain enough information to constrain model design.

explanation: For SEDs that are *explained* by one or more SSRs, one has to document at what level the SSR and SED match. Does the SSR yield a qualitative or quantitative fit for the data?

contradiction: the case where SSRs *contradict* the SED. This situation is subtle, but these subtleties are outside the scope of the abstract.

Another class of entries is that of **predictions.** Some SSRs may not link to SEDs in the literature yet be so interesting that one wishes to recast the result as an empirical Prediction in the hope of encouraging someone to conduct the necessary investigation.

BODB also allows entries for explicit comparison with Related Models.

For brain modeling, BODB provides two ontologies, a "functional ontology" for the Model by listing related Brain Operating Principles (BOPs), and a "structural ontology" in the form of a list of related brain regions. In extending the use of BODB to models of language evolution, new ontologies would be needed, extending the general framework for modeling language evolution described in the companion paper (Arbib 2010).

Clearly, documenting a model in BODB is a non-trivial task. There are three payoffs: (i) The obvious payoff is that researchers will be able to find an analytic account of your model to complement published papers. (ii) A second payoff is the *automatic benchmarking* feature of BODB which allows one to search for related models, then to *automatically* construct a table of the kind shown below:

	Models		
SEDs	***FARS (Fagg - Arbib)***	***Mirror Neuron System (MNS) (Oztop - Arbib)***	***Mirror Neuron System 2 (MNS2) (Bonaiuto - Rosta - Arbib)***
AIP grasp selectivity	Explanation	Support	Support
AIP projection to F5	Support	Support	Support
cIPS projection to AIP	Support	Support	Support
F4 reach selectivity	Scene Setting	Support	Support
F5 canonical visual properties	Support	Support	Support
F5 grasp phase selectivity	Support	Support	Support
F5 grasp selectivity	Support	Support	Support
F5 mirror properties	Scene Setting	Explanation	Explanation
PIP object selectivity	Support	Support	Support
Separation of dorsal and ventral visual streams	Scene Setting	Scene Setting	Scene Setting
VIP location selectivity	Scene Setting	Support	Support
VIP projection to F4	Scene Setting	Support	Support
7b projection to F5		Support	Support
Area 7 is reciprocally connected with STS		Support	Support
Area 7b connection to visual areas		Support	Support
Area 7b connection with AIP		Support	Support
cIPS projection to 7a		Support	Support
F5 mirror - broadly congruent		Scene Setting	Scene Setting
F5 mirror - strictly congruent		Support	Support
F5 mirror - transitive action selectivity		Explanation	Explanation
LIP projection to 7a		Support	Support
MIP projection to 7a		Support	Support

The table lists all the SEDs referred to in the documentation of at least one of the models. It then provides a column for each model showing the relation between the model and the SED. This allows one to rapidly compare the models (the full table for the above comparison indicates that MNS2 explains a number of SEDs not explained by MNS). (iii) This sets the stage for our current work on extending BODB functionality (as distinct from documenting more models).

BODB will become an environment for developing new models. The idea is that one uses a table such as the above (a) to decide whether to extend an existing model or develop a new one, and (b) to choose a set of the SEDs in the table to be used to support the model, or which model development will seek to explain. Thus BODB entry ceases to be an additional chore for model reporting, but becomes a checklist that grows as modeling proceeds, so that a successful BODB entry marks successful completion of the current model.

2. Applying BODB to Models of Language Evolution

The full talk will offer a first pass on new ontologies for indexing models of language evolution, complementing the ontologies of BOPs and Brain Regions currently used to document brain models. Just as one example, we may expect models of *Performance Systems, Developmental Systems, Historical Systems,* and *Evolutionary Systems* to have their own specific ontologies, but the extended BODB must also support the careful analysis of integration across categories, e.g., in showing how choosing an operating principle for developmental systems (e.g., based on Universal Grammar or not) influences the models of historical systems and evolutionary systems.

3. An Invitation

The previous version of BODB is described in a user's manual which is available on-line (Plangprasopchok et al. 2006). During February of 2010, a new release of BODB (based on a major effort by James Bonaiuto) will go live. The talk will include the URL of the new version and outline how researchers on language evolution may begin to enter models into the new BODB.

Acknowledgement: The present research was supported in part by the National Science Foundation Program in Perception, Action, and Cognition under Grant BCS-0924674 (Michael A. Arbib, Principal Investigator).

References

Arbib, M. A. (2010) Modeling language evolution: A general framework with the mirror system hypothesis as a special case. *Workshop on Models of Language Evolution, 8th International conference on the Evolution of Language (Evolang8), Utrecht, The Netherlands,* submitted.

Plangprasopchok, A., Tinroongroj, N., Bonaiuto, J. B. & Arbib, M. A. (2006) User's Manual for the Brain Operation Database: BODB (Version 3). http://nsl.usc.edu/BODB/BODB_Manual__09_05_.pdf.