18

What Can Mathematical, Computational, and Robotic Models Tell Us about the Origins of Syntax?

Herbert Jaeger, Luc Steels, Andrea Baronchelli,
Ted Briscoe, Morten H. Christiansen, Thomas Griffiths,
Gerhard Jäger, Simon Kirby, Natalia L. Komarova,
Peter J. Richerson, and Jochen Triesch

Abstract

This working group staked out the landscape of mathematical and computational modeling approaches to language evolution. The authors attempt both to provide a survey of existing research in this field and to suggest promising novel directions for investigations. The focus was set on cultural evolution; biological evolution is covered in other chapters of this volume. The field is young, manifold, and highly productive, with contributions not only from theoretical linguists but also from theoretical physics, robotics, computational neuroscience, and machine learning. This chapter provides a systematic presentation and comparison of different modeling paradigms, which range from a rigorous mathematical analysis of macroscopic language transmission dynamics to computational (even robotic) studies of complex systems populated by learning agents. Original contributions of modeling research to the understanding of cultural language evolution are highlighted, shedding light on the magnification of learning bias through cultural transmission, restricting the space of possible grammars, coevolution of categories and names, and the emergence of linguistic ontologies.

Introduction

Mathematical and computational models play a crucial role in all sciences and can clearly be helpful to the study of language evolution as well. A model makes certain theoretical assumptions about evolutionary forces and linguistic

representations and then shows what the consequences of these assumptions are on the outcome of language evolution. Of course, the model does not in itself prove that the assumptions are empirically valid but shows whether the assumptions are coherent, have the effect they are believed to have, and are in principle sufficient to generate the phenomena they are claimed to generate. Thus, models provide the opportunity to test different hypotheses about the ingredients that are necessary for languages with certain properties to emerge through processes of transmission and interaction between agents. In turn, they make predictions about the relationship between language acquisition, communicative interaction, and language change that can be assessed through experiments with human participants or robotic agents and through comparison with historical data.

In the field of language evolution, a wide range of models has already been explored, but this is only the beginning. The complexity of the models, the questions they address, and the techniques used to check the validity of current models vary widely. Consequently, the discussions in our group did not and could not be expected to lead to a unified and complete picture, partly because researchers have been looking at entirely different aspects of the enormously complex problem of language evolution and have been using very different methods. Instead we tried to sample the landscape of existing modeling efforts and the representations of grammar and grammatical processing that are used in them. We then surveyed arguments regarding why and how modeling can contribute to the overall language evolution research enterprise, and outlined future research, including possible collaboration with biologists and linguists.

To avoid a possible misunderstanding, we point out that the discussions in our group, and consequently the material in this chapter, focused primarily on the *cultural* evolution of language, to be distinguished from the *biological* evolution, which is in the focus of other contributions in this volume. Nonetheless, investigations of cultural language evolution have implications for research on biological evolution, because if it is found that certain traits of language can be naturally explained by the former, biological mechanisms are relieved from an explanatory load. Conversely, biologically evolved, generic, nonlinguistic information-processing capabilities (e.g., sequential processing mechanisms) yield the scaffolding for cultural evolution.

Another preparatory remark is that modeling efforts adopt approaches that are quite standard in other domains of complex systems science, but may be relatively new to linguists. For example, there is often an effort to seek simplified models to clearly pin down the assumptions and, in many cases, to make the models tractable from a mathematical point of view. Modelers typically focus on replicating statistical distributions of language phenomena rather than matching directly the particulars of a given human language. They will first consider communication systems that have only a rudimentary resemblance to language before increasing the complexity further, step by step. Alternatively, they will make assumptions about certain aspects of language interaction (such

as joint attention or perception) to make simulations viable at all. Some models are not about language per se but address the preconditions for language, such as cooperation. It is therefore important to keep in mind that the modeling work discussed here is primarily concerned with investigating the consequences of hypotheses rather than trying to model in detail and in a realistic way the origins and evolution of human language.

Paradigms for Studying Language Evolution

Discussions in our group arose from the multifaceted experience of the participants with computer-based simulations of language dynamics, robotic experiments, and mathematical analysis. We are not aware of any generally accepted way of characterizing or classifying computational modeling approaches in the natural or social sciences. In the present context, we could nevertheless identify a number of different modeling paradigms that have grown up historically based on the shared interests of the researchers involved in exploring them. Each paradigm frames the process of language evolution in a particular way, focuses on some of the forces that might play a role, and then examines specific fundamental questions through concrete models and experiments. Within each paradigm we have seen the development of mathematical models, computational or robotic experiments, and psychological experiments with human subjects. Of course, the distinctions between paradigms made here are to some extent arbitrary and not always clear-cut. There are continuous dimensions linking these paradigms and hence considerable opportunities for cross fertilization. Moreover, we anticipate that additional modeling paradigms may spring up in the future to explore other aspects of the vast research domain of language evolution.

A first distinction that can be made is between agent-based models, which try to pin down the cognitive and social processes that could give rise to forms of language, and macroscopic models, which aggregate the behavior of a population and then formulate equations defining the evolution over time among these aggregate quantities. Another dimension for categorizing the models concerns the importance given to cultural transmission, cognition, or biology, which has given rise to iterated learning models, language games, and genetic evolution models. In Figure 18.1 we illustrate schematically two main dimensions on which the paradigms differ.

Agent-based Models

Agent-based models center on models of individual language users as members of populations. The agents are given certain cognitive capabilities (e.g., a particular learning strategy) and made to interact (e.g., in the simulation of a teacher–learner situation or a communicative interaction between two
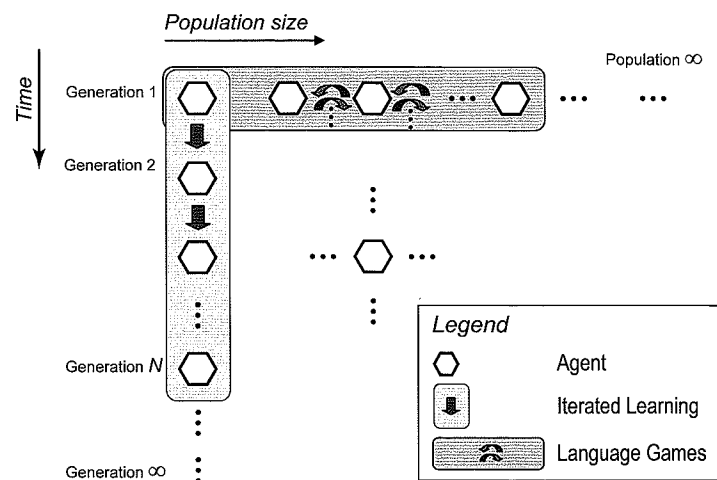
Figure 18.1   Schematic "coordinate system" comparing agent-based language evolution paradigms. The simplest models within the iterated learning paradigm focus on transmission across *generations* of agents in a singleton chain of teacher–learner dyades; language games focus on how language constructs emerge and evolve in interactions between agents. Numerous other paradigms can be seen as mixtures and ramifications of these two.

individuals). By simulating the effect of a large number of interactions, agent-based models can study under what conditions language systems with similar properties as human natural languages can appear. Agent models vary greatly in complexity, ranging from simple statistical "bag of words" language models to robots using complex grammatical and semantical representation formalisms to communicate with each other in a dynamical environment.

Three types of agent-based models have been developed: iterated learning models which focus on understanding the role of cultural transmission, language game models which emphasize the role of communication and cognition, and genetic models which explore the role of biological evolution.

Iterated Learning

The first paradigm, which has already been explored quite extensively, is known as the *iterated learning paradigm*. It focuses on understanding the relationship between properties of the individual and the resulting structure of language by embedding a model of an individual learner in a so-called "transmission chain" (also known as "diffusion chain"; for further details see Kirby et al. and Briscoe, both this volume; for a review of this approach to studying cultural evolution more generally, see Mesoudi 2007). In these models, the linguistic behavior of one individual becomes the learning experience of another individual, who in turn goes on to produce behavior that will be input for a

third individual and so forth. The focus of this framework is on the contribution of learning in shaping the process of cultural transmission, with the goal of specifying precisely the relationship between constraints and biases provided by biology and the universal properties of linguistic structure. The idea is that a fundamental challenge for language is to be repeatedly transmitted between individuals over generations, and the transmission process is imperfect in important ways (e.g., learners have particular biases, they only see a subset of the language, there is noise in the world). The result is an adaptive system whereby language evolves culturally in such a way so as to give the appearance of being designed for transmission fidelity.

The main simplification in many (but not all) of the models of this "iterated learning" process is that the transmission chain consists of a single individual at each generation and involves only vertical transmission (i.e., transmission between generations). This simplification allows researchers to focus on the sole contribution of the learning bias plus the nature of the selection of training data (e.g., number of examples), although it leaves out many of the factors associated with horizontal transmission (e.g., selection of models to learn from, having shared communicative goals, and population structure). One avenue for future research is to explore the implications of other, more realistic models of populations, while maintaining the emphasis on the role of transmission in shaping language structure. For a recent review of general cultural evolution models, see McElreath and Henrich (2008).

Examples of iterated learning models are presented by Kirby et al. (this volume). An emphasis in many of these models has been the explanation of the emergence of compositional structure in language. Compositionality, along with recursion, is the fundamental feature of human syntax that gives us open-ended expressivity. It is also arguably absent in any other species, despite the prevalence of communication in nature. Accordingly, it is an important target for explanation. Using mathematical, computational, and experimental models, researchers have examined the conditions under which compositionality and the relationship between compositionality and frequency may emerge. Specifically, these models suggest that compositionality arises when there is a "bottleneck" on the cultural transmission of language; in other words, where learning data is sparse.

Language Games

The second class of models investigates the role of embodiment, communication, cognition, and social interaction in the formation of language. Instead of modeling only teacher–learner situations, as in iterated learning approach, it models the communicative interactions themselves in the form of language games. A language game is a situated embodied interaction between two individuals within a shared world that involves some form of symbolic communication. For example, the speaker asks for "a cup of coffee" and the hearer gives

it to her. When speaker and hearer have shared conventions for solving a particular communicative problem, they use their existing inventory in a routine way. When this is not the case, the speaker requires the necessary cognitive capabilities to extend his inventory (e.g., expanding the meaning of a word or coercing an existing word into a new grammatical role), and the hearer requires the ability to infer meanings and functions of unknown items, thereby expanding his knowledge of the speaker's inventory.

In typical language game models, the individuals playing language games are always considered to be members of a population. They interact only in pairs without any centralized control or direct meaning transfer. There is unavoidable variation in the population because of different histories of interaction with the world and others; however, proper selectionist dynamics, implemented by choosing the right alignment and credit assignment strategies for each individual, causes certain variants to be preferred over others. Language game models often operate with a fixed population because they examine the thesis that language emerges and evolves by the invention, adoption, and alignment strategies of individuals in embodied communicative interactions. In addition, many experiments have been done in which a flow is organized in the population with members leaving or entering the population, thus demonstrating that the model handles cultural evolution as well.

By now there have been dozens of experiments in language games that explore how different aspects of language may arise (see Steels, this volume). The simplest and earliest game studied is the *naming game*, in which agents draw attention to individual objects in the world by using (proper) names (Steels 1995). Guessing games have been used to study the coevolution of perceptually grounded categories and words (Steels and Belpaeme 2005), flexible word meanings (Wellens et al. 2008), and the emergence of spatial language (Steels and Loetzsch 2008). Description games have been used in experiments on the emergence of grammar, in particular, case grammar (Van Trijp 2008).

Language games have been explored further from three angles: through mathematical analysis, particularly using the methods of statistical physics; through computational simulations and robotic experiments; and through experiments with human subjects as carried out, for example, by Galantucci (2005) and Pickering and Garrod (2004). Robotic experiments are particularly useful if one wants to study the question of how embodiment plays a role in language evolution. Data on actual language change, which comes from historical linguistics and socio-linguistics, are currently being used to constrain the repair and consolidation strategies of agents in grammatical language games; data from cognitive linguistics and cognitive semantics, in particular, are used to constrain the range of possible conceptualizations that could be the target of experiments. The theoretical tools developed in statistical physics and complex systems science have recently acquired a central role for the study of language games. The suite of methods developed in these fields has indeed allowed us to address quantitatively such issues as the scaling of relevant features of the

models with the system size (e.g., convergence time or memory requirements; Baronchelli et al. 2006b, 2008), the impact of different underlying topology on global behaviors (e.g., homogeneous mixing [Baronchelli et al. 2008] vs. regular lattices [Baronchelli et al. 2006a] vs. complex networks [Dall'Asta et al. 2006a, b]), and the detailed study of convergence dynamics (Baronchelli et al. 2008). Thus, for example, it has been shown that complex networks are able to yield, at the same time, the fast convergence observed in unstructured populations and the finite memory requirements of low dimensional lattices (Dall'Asta et al. 2006a, b). Moreover, agents' architectures and interaction rules have been significantly simplified to allow thorough analysis, and this has enabled us to pinpoint the crucial ingredients responsible for the desired global coordination. The pursuit of simplicity, along with the novelty of the complex systems approach to this field, has thus far mostly limited the investigations to the study of the naming game and category game (in which the population ends up with a shared repertoire of categories) (Puglisi et al. 2008). Current research is trying to tackle higher-order problems, such as the emergence of compositionality (de Beule 2008). Experiments with human subjects show that humans can evolve communication systems, although some are better than others, mostly because of differences in social attitudes. Of course the greatest challenge is to scale these experiments up to the level of grammatical languages. Recent examples already showing the formation of case grammars, tense–aspect–mood systems, or determiner systems give reason for optimism (see, e.g., Van Trijp 2008).

Genetic Evolution

A third class of models explores the role of biology by modeling the genetic transmission of language. Agents are created based on a model of a genome that directly codes the lexicon or grammar of their language. Agents then engage in interactions that determine their fitness, and based on communicative success they have a higher chance of reproducing in the next generation. Due to random mutations and crossover, offspring have slightly different genomes, possibly giving higher communicative fitness which then leads to further propagation. These models use very similar techniques to those used in genetic algorithms, and they sprang up first in the context of artificial life (Cangelosi and Parisi 1998). Given that the explicit genetic coding of lexicon and grammar is highly implausible from a biological point of view, more recent models have considerably weakened this assumption and encode only strong biases and universal constraints on possible languages. This is particularly the case for the ENGA model (Szathmáry et al. 2007). ENGA is an ambitious framework that covers not only the genetics but also the neurodevelopmental processes in a biologically realistic way. Linguistic inventories are not coded genetically; they are acquired through a learning process. The ENGA model,

therefore, attempts to cover the entire spectrum: from genetic to developmental and learning processes.

Aggregate Models

In addition to agent-based models, extensive research has constructed macroscopic models of language evolution and language dynamics.

Game Theoretic Models of Language Evolution

The main paradigm being explored here draws from the tradition of evolutionary game theory to focus on the role of imitation in cultural transmission. Imitation (or reuse) applies both to the adaptation of linguistic performance between adult speakers and the acquisition of language by infants. Imitation is framed as a form of replication. An evolutionary dynamics ensues in any population of replicating entities, provided the entities in the population vary in certain heritable characteristics, and replicative success is correlated with this variation. This is a crucial difference to the iterated learning paradigm, where every individual grammar participates equally in language replication. However, the game theoretic model—as a form of a selectionist model—assumes faithful replication, while replication under iterated learning may be imperfect. Under certain simplifying assumptions—like the postulation of an infinite population and a continuous time—such evolutionary dynamics can be described by a system of ordinary differential equations. In language evolution this dynamics is necessarily nonlinear because selection is frequency dependent. This can, for instance, be illustrated by the development of vocabulary: whether a candidate for a neologism catches on in a linguistic community (i.e., becomes replicated) depends on whether or not there already is another word for the same concept within this linguistic community. This indicates that the overall frequency distribution of words is a decisive factor for the fitness of each individual word. A similar point can be made for other linguistic units, ranging from phonemes to syntactic constructions.

Frequency-dependent selection can be modeled by means of replicator dynamics within the mathematical framework of evolutionary game theory (Maynard Smith 1982; Hofbauer and Sigmund 1998). In its simplest form, a model of a communication game consists of:

- a space of meanings and a space of forms,
- a space of production grammars (mappings from meanings to forms),
- a space of comprehension grammars (mappings from forms to meanings), and
- a utility function (i.e., a measure of success for a pairing of grammars, depending on the success of communication and complexity of the grammars involved).

Further parameters may be added, like a biased a priori probability distribution over meanings, or a confusion matrix for noisy transmissions of forms.

There are several off-the-shelf theorems from biomathematics regarding stability conditions for evolutionary games. Such theorems sometimes render it straightforward to identify the attractor states of the replicator dynamics without actually delving into the complexities of the underlying nonlinear differential equations.

The biomathematics literature contains a variety of results concerning the evolution of communication, where strategies ("grammars") are assumed to be innate and replication is interpreted in the biological sense (e.g., Wärneryd 1993; Trapa and Nowak 2000; Nowak and Krakauer 1999; Nowak et al. 1999; Jäger 2008a). These authors mainly consider biological evolution, and they assume that communicative success is correlated with biological fitness (i.e., the number of fertile offspring). However, their results are general enough that they can be extrapolated to cultural evolution. The background assumption here is that communicative success of a certain behavioral trait is positively correlated with its likelihood to be imitated (i.e., its cultural fitness). Possible applications of evolutionary game theory to the study of the cultural evolution of language (in the sense described above) have been investigated (Jäger 2007, 2008b; Jäger and van Rooij 2007).

Game theoretic research in language evolution has suggested a formal framework which is quite useful within this paradigm. Universal Grammar or a preexisting bias of grammar learning can be represented in the following abstract manner. Suppose we have a finite alphabet (i.e., a finite set of symbols) (Nowak et al. 2001; Komarova et al. 2001; Komarova and Nowak 2001, 2003; Nowak and Komarova 2001). A language is a probability distribution defined on a set of strings composed of the symbols of the alphabet. The allowed languages can be represented as probability distributions on a collection of (intersecting) sets. Then a learning mechanism is a way to "navigate" in this collection of sets. Pair-wise similarity among languages can be expressed as a matrix. The process of learning is thus a sequence of hypotheses of a learner in response to the input of a teacher (or teachers), which is a number of strings compatible with the teacher(s)' grammar. This framework allows one to use the machinery from mathematical learning theory and connect natural language evolution with insights from computer science/machine learning.

Summary

There are obvious relations, complementarities, and continuities between these approaches and paradigms. The game theoretic paradigm focuses on the selectionist dynamics of the language itself, whereas language game models use an agent-based approach, focusing on the cognitive mechanisms by which agents use, invent, and coordinate language so that the selectionist dynamics of language emerges. The iterated learning paradigm focuses on the role of bias and

the vertical transmission bottleneck and tends, therefore, not to integrate the issue of communicative success, cognitive effort, or population dynamics into the models. The language game paradigm considers vertical transmission as an additional but not crucial effect on language evolution. Pursuing these different approaches provides us with the opportunity to explore how different factors (e.g., learning, communication, and population structure) influence the process of language evolution.

Linguistic Representations and Processes

Given that syntax was the focus of this Forum, it is relevant to examine what kind of representations for grammar are being used in language evolution models and what kind of syntactic operations and grammatical processes have been incorporated into these models. Researchers working on iterated learning and game theoretic approaches generally try to use existing "symbolic" formalisms or neural network models. Some have argued, however, that the requirements of evolvability put additional constraints on the nature of grammatical representations and processing, and this has led to some work on novel grammar formalisms which can cope with emergent grammar.

Symbolic Grammars

There are a variety of grammatical formalisms in the theoretical linguistics literature, some of which have been utilized in evolutionary models whereas others, such as Minimalism (Chomsky 1995), have not (possibly because they are less easily embedded in theories of processing). Examples of formalisms that have been deployed with minimal modification include optimality theory (Jäger 2004), extended categorial grammar (Briscoe 2000), and context-free grammars (Zuidema 2002). All such models require the embedding of the formalism into a theory of grammar learning and processing. Modelers have drawn on existing proposals from the literature, such as Bayesian parameter estimation, compression-based algorithms, or nonstatistical parameter-setting algorithms for implementing the learning mechanisms used in vertical transmission (Griffiths and Kalish 2007; Briscoe 2000).

Simple Recurrent Networks

Other language evolution models have avoided the explicit representation of hierarchical structures, syntactic and semantic categories, and grammatical rules, deploying distributed and subsymbolic representation. A popular alternative is simple recurrent networks (SRNs; Elman 1990). In SRNs, knowledge of language is learned from the presentation of multiple examples from which the networks learn to process syntactic structure. The general aim of

such models is to capture observable language performance, rather than idealized linguistic competence (Christiansen 1992; Christiansen and Chater 1999). Much of this work has an emphasis on the integration of multiple sources of probabilistic information available in the input to the learner/speaker/hearer (e.g., from the perceptuo-motor system, cognition, socio-pragmatics, and thought as discussed by Kirby et al., this volume). Although much of this work tends to target small fragments of language for the purpose of close modeling of psycholinguistic results (e.g., Christiansen and Chater 1999; MacDonald and Christiansen 2002), some efforts have gone into scaling up models to deal with more realistic language samples, such as full-blown child-directed speech (Reali et al. 2003). In this framework, grammatical processing can be conceptualized as a trajectory through a high-dimensional state-space afforded by the hidden unit activations of the network (e.g., Elman 1990), potentially suggesting an alternative perspective on constituency and recursion in language (Christiansen and Chater 2003). These models do not include explicit grammar formalisms, but the behavior of the networks can in some cases be described in terms of such formalisms.

Formalisms Designed for Grammar Evolution

Some researchers have developed novel formalisms to be used specifically in language game experiments. This is particularly the case for *fluid construction grammar* (FCG). FCG (de Beule and Steels 2005) uses representational mechanisms already employed in several existing symbolic grammar formalisms like *head-driven phrase structure grammar* (HPSG; Sag et al. 2006) or lexical–functional grammar (Kaplan and Bresnan 1982) such as a feature/ structure-based representation of intermediary structures during parsing and production, a constraint-based representation of linguistic rules so that they can be applied in a bidirectional fashion, and unification-style mechanisms for the application of these rules. FCG is in line with other construction grammar formalisms, such as embodied construction grammar (Bergen and Chang 2004), in the sense of supporting the explicit representation and processing of constructions, which is de-emphasized in Minimalism. However, FCG has various additional facilities to enable language evolution experiments:

1. Individual agents represent a multitude of hypotheses about the emerging language and are therefore able to handle variation in language use.
2. Rule application is flexible allowing the violation of constraints and robust parsing and production so that sentences can be understood even if they are not entirely grammatical (according to the preferred grammar of the agent).
3. The different variants compete within the individual when it has to make decisions about how to express something or interpret something

and, as an emergent effect, within the population for dominance in the emergent language.

4. Rather than coding systematicity in terms of more abstract rules, FCG maintains links between the rules, based on how the rules are formed through composition of other rules.

These links are then used for assigning credit or blame after a game, allowing the implementation of a multilevel selectionist dynamics. Because of these features, FCG exhibits dynamical systems properties seen in network representation systems of grammar which do not rely on symbolic structures, such as connectionist networks or recurrent neural networks (RNNs), while at the same time incorporating many ideas from decades of research into theoretical and computational linguistics.

Summary

The computer simulations carried out in the evolution of language research rest on a variety of formalisms to represent the inventory of the lexicon and grammar of the emerging language in the first place. In choosing a particular formalism, the researcher makes a commitment to what aspects of language are isolated for an inspection of their role in evolutionary dynamics, and what others are (implicitly) excluded.

How Can Modeling Inform the Study of Language Evolution?

Although mathematical and computational modeling of language evolution is still in its infancy, there are already quite a few results that show the power of the approach, and which may be of interest to biologists and linguistics. Computational modeling, like in any other field, enables two powerful avenues for accruing scientific insight:

1. *Formal analysis.* Computational models have to be rigorously formalized to make them operational on computers. When a simulation is running, all aspects of the simulation can be recorded, including the population aspects. The same is true in robotic experiments where all perceptual states, motor states, and the full details of all processes going into language production and understanding can be tracked; this is not possible with human subjects. This full access to relevant data makes the models amenable to mathematical analyses. Typical questions that can be answered by the analytical methods provided by nonlinear dynamics, game theory, and statistical physics concern asymptotic properties of evolutionary dynamics, the dependence of these dynamics on scaling parameters, or the prediction of sudden and dramatic changes (phase transitions).

2. *Simulation studies.* Carrying out simulations on a computer differs from carrying out real-life experiments in two crucial ways: (a) the simulated piece of reality is *completely* specified, and (b) one has *full control* over varying experimental conditions. There are risks and opportunities under these circumstances. An obvious pitfall is that the simulation may miss a crucial component of the real-life target system—this is the problem of abstraction. However, it should be noted that, in principle, the same problem is present in experimental designs involving human subjects: a particular experimental design may prevent real-life-relevant mechanisms from taking effect. The benefits added to empirical experiments (which remain indispensible) by simulation studies are, in our view, the following:

-   A systematic *exploration of large hypothesis spaces* is made possible due to the speed and low cost of simulations. This facilitates both the generation of new scientific hypotheses and the testing of existing ones.
-   Model simulations can give existence proofs for the efficiency of certain mechanisms to achieve a certain effect—always, of course, modulo the modeling assumptions.
-   In a related vein, model simulations can give nonuniqueness proofs if the same ultimate effect can be obtained by different mechanisms. Such demonstrations are helpful in precluding an early "contraction" to a single explanatory venue in theory development.
-   Simulations are replicable across different laboratories by sharing code.
-   Critiquing and improving simulation setups is transparent, because it is explicit how assumptions become operationalized in the designs.

If one is carefully conscious of the assumptions that go into a simulation model, research based on such models can decidedly "open up" the space of possible theories in a field, raising the awareness of alternative theories. To demonstrate this point, we present a number of examples that have arisen from ongoing work.

Magnification of Learning Bias through Cultural Transmission

Mathematical analyses of the iterated learning model provides some interesting insights into the relationship between the inductive biases of language learners—the factors that lead them to find it easier to learn one language than another, as might be the result of genetic constraints on language learning—and the kinds of languages that will be spoken in a community. As discussed by Kirby et al. (this volume) and Briscoe (this volume), one way to capture the inductive biases of learners is to assume that they identify a language from a set of utterances by applying Bayesian inference, with a "prior" distribution

encoding which languages learners consider more probable before seeing any data. Languages with higher prior probability can be learned from less evidence, and the prior thus reflects the inductive biases of the learner. Analyses of iterated learning with Bayesian agents show that the relationship between the prior and the languages that are ultimately produced via cultural transmission can be complex (Griffiths and Kalish 2007; Kirby et al. 2007). Specifically, iterated learning can magnify weak inductive biases, with a slight difference in the prior probabilities of two languages resulting in a significant difference in the probability of those languages being produced via cultural transmission. These mathematical results suggest that strong genetically encoded constraints on learning may not be necessary to explain the structure of human languages, with cultural evolution taking on part of the role that might otherwise have been played by biological evolution.

Restricting the Space of Possible Grammars

It is tempting to reconstruct the notion of a linguistic universal as a property that every language with a grammar that can be cognitively represented and learned by humans (i.e., a language that conforms to "Universal Grammar" in the Chomskyan sense) shares. Evolutionary models indicate that there may be other sources of universals. Briefly put, a *possible language* must also be attainable under the evolutionary dynamics of language transmission.

In Jäger (2004), this basic idea is illustrated with a particular implementation. According to optimality theory, Universal Grammar defines a finite set of constraints, and each particular grammar is characterized by a linear ordering of these constraints. To account for certain strong typological tendencies, Aissen (2003) proposed restricting the space of possible grammars further by imposing certain subhierarchies of constraints that are never violated.

Following proposals by Boersma (1998), Jäger implements a stochastic learning algorithm for optimality theoretic grammars. However, unlike Boersma, Jäger assumes that language acquisition is bidirectional (i.e., the learner tries both to mimic the production behavior and the comprehension behavior of the teacher). It turned out that some constraint rankings are strictly not learnable at all. Among the remaining space of learnable grammars, some are more robustly learnable than others. After iterating the learning procedure a few dozen or hundred of times (where in each generation, the former learner becomes the teacher and produces utterances on the basis of his acquired grammar), only constraint rankings that conform to Aissen's prediction were observed.

The Coevolution of Categories and Names

One of the big debates in language studies concerns the question of how far perceptually grounded categories (e.g., such as colors) influence and are influenced by language that expresses these categories. From a Whorfian point of

view there is a strong interaction, whereas those arguing for strong modularity argue that categories are innate or induced from empirical data and language just labels existing categories. Although color categorization and color naming does not relate directly to grammar, we include this theme here because it exemplifies the quality of insight that can be obtained from modeling studies, and because categorization and naming are prerequisites for grammatical language. Research on the use of language games for studying the coevolution of categories and names began with Steels and Belpaeme (2004); here, agent-based models of color naming and categorization were developed and systematically compared. This work showed that although a genetic evolution of color categories was possible, it not only took a long time, but also did not lead to a system that was adaptive, and it certainly did not lead to universal categories unless populations remained homogeneous. It also demonstrated that a purely learning-based approach did not lead to an explanation for trends in color categories nor to sufficient coherence in a population to explain how a successful communication was possible. More recently, this research has been extended in two directions.

The category game (Puglisi et al. 2008) is a language game that aims at describing how a population of agents can bootstrap a shared repertoire of linguistic categories out of pairwise interactions and without any central coordination. The main result is the emergence of a shared linguistic layer in which perceptual categories are grouped together into emerging linguistic categories to guarantee communicative success. Indeed, while perceptual categories are poorly aligned between individuals, the boundaries of the linguistic categories emerge as a self-organized property of the whole population and are therefore almost perfectly harmonized at a global level. Interestingly, the model reproduces a typical feature of natural languages: despite a very high resolution power and large population sizes, the number of linguistic categories is finite and small. Moreover, a population of individuals reacts to a given environment by refining the linguistic partitioning of the most stimulated regions, while nonuniform JNDs (e.g., the human JND function relative to hue perception) constrain to some extent the structure of the emergent ontology of linguistic categories.

Another simple framework has been designed to investigate the influence of various realistic features (linguistic, psychological, and physiological) on the shared color categorization (Komarova and Jameson 2008). As a result of a number of iterations of the appropriate game, a population of agents arrives at a shared categorization system, which possesses the following qualities: (a) the exemplar space is equipartitioned into a (predictable) number of distinct, deterministic color categories, (b) the size of color categories is uniquely defined by the pragmatic similarity parameter, and (c) the location of category boundaries possesses rotational symmetry. Empirical data of confusion spectra of abnormal color observers can be incorporated to generate specific color boundary predictions and to deduce how the color categorization of various

populations is influenced by the population inhomogeneities (Komarova and Jameson 2008).

The Emergence of Linguistic Ontologies

The final example, which shows how modeling can lead to the opening up of new theoretical avenues and ideas in language evolution, comes from the domain of grammar. Grammar exploits syntactic devices (e.g., word order or morphology) to express additional aspects of meaning, such as discourse structure, thematic relations (predicate–argument structure), tense–aspect–mood, determination, scoping constraints on anaphora. In all current linguistic theories, the rules of grammar are expressed using an ontology of syntactic and semantic categories. These syntactic categories include, for example, parts of speech (e.g., noun, verb, adverb), types of constituents (e.g., noun phrase, relative clause), syntactic constraints (e.g., agreement, precedence), and syntactic features (e.g., nominative, masculine, neuter). The semantic categories include classifications of temporal aspects in terms of tense, aspect, or mood, semantic roles such as agent or beneficiary, categories used for conceptualizing discourse, like topic/comment, different shades of determination (e.g., definite/indefinite, count/mass), classifiers (as used in Bantu languages), deictic references both for use inside and outside discourse, epistemtic distinctions, and so on. A complex grammar undoubtedly requires hundreds of such categories. Thus, to understand the origins and evolution of language, we must know the origin of such a linguistic ontology.

There is a common (usually hidden) assumption among many theorists that linguistic ontology is universal and innate, but that does not explain how it originates. Typologists have argued that linguistic categories are to a large extent language dependent (Haspelmath 2007), and historical linguists have shown that categories change over time (Heine and Kuteva 2008). This suggests that linguistic categories may be similar to categories in other domains of cognition (e.g., the color categories), in the sense that they are culturally constructed and coordinated.

Recent language game experiments in the formation of a case grammar (Steels, this volume) have shown that the formation of linguistic ontologies is entirely possible. Concretely, semantic roles as needed in case grammar have been shown to arise when agents are trying to reuse by analogy semantic frames that have already been expressed in the emergent language. This reuse becomes licensed when particular predicate–argument relations are categorized in the same way as those already used in the existing semantic frames. Progressively, semantic roles get established and refined, partially driven by the semantic analogies that make sense in the real-world domain that generates the topics in the language game and partly by the conventions that are being enforced by the emergent language (Van Trijp 2008).

Summary

The examples discussed here illustrate some of the ways in which models of the cultural evolution of language can contribute to our understanding of its origins. By identifying what aspects of the properties of languages can be produced by cultural evolution alone, these models remove some of the explanatory burden from biological evolution, providing a more realistic target for research into the origins of language. In broad terms, these models illustrate how learning, communication, and population structure affect the languages that emerge from cultural evolution, providing potential explanations for two of the most important aspects of human languages: their consistent properties across communities—language universals—and the coherence of linguistic systems within communities. In iterated learning models, universals emerge as the result of learning biases or the goals of communication, and coherence is the result of the strength of these biases and the structure of the interactions with other individuals. In language game experiments, universal trends emerge due to constraints coming from embodiment, the cognitive mechanisms recruited for language, the challenge of communication, and the selectionist dynamics that emerges in populations of adaptive communicating agents. While there are still many questions to explore, these basic results help to illustrate the kinds of forces that influence the structure of human languages.

Suggestions for Future Research

Given that there is a broad variety of paradigms and modeling efforts, there are also many possible avenues for deepening current results or for exploring new avenues of research. Here we describe a number of suggestions without any claim to be exhaustive. In general, we can expect models to be developed that focus on quite different aspects of language evolution and that will be formulated at very different levels of abstraction. It will be important to establish the relationships between these models, such as identifying to what extent a simpler and more abstract model can be understood as an approximation to a more elaborate one.

Tighter Coupling between Models and Laboratory Experiments

It is important for future research to develop a tighter coupling between models and laboratory experiments. Currently, there are two ways in which conducting laboratory experiments in cultural evolution can complement the insights provided by mathematical and computational models. First, they provide a direct way of testing the predictions of these models, allowing us to ensure that the claims that we make about cultural evolution are actually borne out when these processes involve real people rather than abstract agents. For example,

Kalish et al. (2007) and Griffiths et al. (2008) have conducted direct tests of the key prediction that arises from models of iterated learning with Bayesian agents (i.e., structures that are easier to learn will be favored by the process of cultural transmission) by conducting laboratory experiments in which the structures transmitted by iterated learning were categories and functional relationships between variables for which previous research in cognitive psychology had established results on the difficulty of learning. Laboratory experiments, however, can also be valuable for a second reason: they provide a closer approximation to the true processes involved in language evolution. The models discussed earlier make assumptions both about how information is passed between agents and the learning mechanisms used by those agents. Conducting laboratory experiments in which information is passed between agents in the way described by a model, but where the agents are real human beings, removes one level of approximation from these models, allowing us to explore the plausibility of processes of cultural transmission as an account of why languages have the properties they do (Dowman et al. 2008). The experiment described by Kirby et al. (this volume) is of this kind, showing that iterated learning with human learners produces compositional structures. Further experiments that test models of language evolution and evaluate the impact of different forms of cultural transmission can help us develop models that provide a closer match to human behavior and assess the contributions of different kinds of evolutionary forces.

Tighter Coupling between Models and Data from Historical Linguistics

Much is known about the historical evolution of human languages over the past 5000 years. This research shows that there are recurrent patterns of grammaticalization and lexical change, and detailed case studies exist to explain, for example, how a language has developed determiners, a case system, or a tonal system (see, e.g., Heine and Kuteva 2008). It is therefore obvious that these results should constrain models of language evolution. Although it will never be possible to reconstruct the actual evolution of human languages, it might be possible to see similar grammaticalization phenomena as in human languages.

Modeling the Potential Role of Exaptation on Language Evolution

It is widely assumed that language in some form or another originated by piggybacking on a preexisting mechanism—exaptations—not dedicated to language. A possible avenue of language evolution modeling involves testing the possible effects for language evolution of particular hypothesized exaptations. For example, improved sequential learning of hierarchically organized structure in the human lineage has been proposed as a possible preadaptation for language (Christiansen and Chater 2008; Conway and Christiansen 2001), in

part based on work in language acquisition (Gómez and Gerken 2000) and genetic data regarding the potential role of *FOXP2* in sequential learning (Hilliard and White, this volume). Reali and Christiansen (2009) have explored the implications of such assumptions by determining the effect of constraints derived from an earlier evolved mechanism for sequential learning on the interaction between biological and linguistic adaptation across generations of language learners. SRNs were initially allowed to evolve "biologically" to improve their sequential learning abilities, after which language was introduced into the population, comparing the relative contribution of biological and linguistic adaptation by allowing both networks and language to change over time. Reali and Christiansen's (2009) simulation results supported two main conclusions: First, over generations, a consistent head-ordering emerged due to linguistic adaptation. This is consistent with previous studies which suggest that some apparently arbitrary aspects of linguistic structure may arise from cognitive constraints on sequential learning. Second, when networks were selected to maintain a good level of performance on the sequential learning task, language learnability is significantly improved by linguistic adaptation but not by biological adaptation. Indeed, the pressure toward maintaining a high level of sequential learning performance prevented biological assimilation of linguistic-specific knowledge from occurring. Similarly, it may be possible to investigate the potential effects of other hypothesized exaptations on the relative contribution of cultural evolution and genetic assimilation to language evolution.

Along the same line, several language game experiments have examined how generic cognitive mechanisms could become recruited for language, pushed by the needs to solve specific problems in communication or in bootstrapping an efficient system (Steels 2007). For example, perspective and perspective reversal is often lexicalized in human languages to avoid ambiguity from which point of view a spatial relation should be interpreted.

Effects of Biased Unfaithful Copying

When empirical predictions are derived from dynamical models, the notion of an *equilibrium* is central. In the evolutionary context, we expect systems to spend most of their time in an *evolutionarily stable state*. The insights from historical linguistics, especially regarding *grammaticalization*, indicate that language never actually reaches such a stable state.[1] Rather, languages perpetually change in a partially predictable way. Complex morphology tends to be reduced over time and to disappear altogether eventually. An example is the loss of case distinctions from Latin (five cases) to French (no case distinctions). On the other hand, lexical morphemes are recruited to serve grammatical

---

[1] This statement might be too bold in its generality. Some aspects of language are certainly in equilibrium most of the time. A good example might be vowel systems.

functions. A recent example is the use of the item "going to" in contemporary English to express future. This recruitment usually concurs with phonological reduction, like the change from *going to* to *gonna*. Grammatical words tend to get further reduced to affixes; one example would be the regular German past tense morpheme *t* that is originally derived from the Germanic verb *tun* (in English: *do*).

The macroscopic consequence of these processes is that languages continually change their grammatical type, moving from synthetic to analytic, due to reduction of morphology, and back to synthetic, due to recruitment of lexical items for grammatical purposes and their subsequent reduction to affixes. The underlying microdynamics involves biased unfaithful copying (i.e., words and phrases are not imitated verbatim but phonetically reduced and semantically modified). The challenge for evolutionary models is to connect these two aspects in such a way that the directedness of language change is connected to empirical insight about unfaithful replication in language use. Deutscher (2005) has proposed a verbal model that resembles the socio-linguistic arguments of Labov (2001). Individuals often innovate new speech forms in an effort to find a more emphatic or colorful way of phrasing an idea or grammatical function. Conventional forms bore us while prose or speech stylists who play with the limits of convention attract attention. When prestigious people do this, the new speech form tends to spread. Sometimes the motivations for innovation are social; people seem to favor forms of speaking that differentiate them from social others. In other words, linguistic equilibria are weakly constrained in that communicating individuals must agree *sufficiently* on meanings for communication to be possible. Yet a speech community can easily cope with a modest rate of innovation driven by social and aesthetic forces. To our knowledge, these mechanisms have not been incorporated into formal models except in the special case of symbolic markers of group boundaries (McElreath et al. 2003).

Long-term Language Change Dynamics: A Mathematical Perspective

It appears that language modeling poses challenges for existing mathematical methods commonly used to describe emerging and dynamical real-life phenomenon. A ready example comes from language games. Language game solutions may vary with regards to their stability properties, depending on the type/purpose of the model use and on the exact question we address. In certain situations, interesting quasi-stable solutions are attained. One instantiation comes from modeling color categorization in people, where the shared population categorization solution cannot be described as a stable solution of a dynamical system, or a stationary probability distribution of a stochastic process. In the category game (Puglisi et al. 2008), even though the only absorbing state is the trivial one in which all the agents share the same unique word for all their perceptual categories, there are clear signatures of a saturation with time

of metastable states with a finite and "small" number of linguistic categories. This observation suggests an analogy with glassy systems in physics (Mezard et al. 1987) and is also confirmed by quantitative observations. Thus, in this framework, interesting solutions would be long-lived (strongly, e.g., exponentially, dependent on the population size) pre-asymptotic states. In other models of color categorization, the shared population categorization solution appears dynamically stable on a certain timescale, but it may drift or cycle (while retaining global topological structure) on longer timescales, depending on the particular constraints. Mathematical properties of such solutions have not been investigated in detail but their understanding is important because conventional methods do not grasp the relevant properties of such solutions. The application of new mathematical technologies thus developed will be wide, as it has implications in the dynamics of populations of learners trying to achieve shared solutions on (possibly very complex) topological semantic spaces.

Selectively Neutral Mechanisms of Linguistic Evolution

A further direction for future research is to understand to what extent processes of selection are necessary to explain the properties of languages. In biology, selectively neutral processes such as mutation and genetic drift have been identified as playing a significant role in accounting for genetic variation (Kimura 1983). It remains to be seen whether linguistic variation is best analyzed as the result of selective pressures acting on the properties of languages, or the outcome of selectively neutral processes that are the cultural equivalents of mutation and drift. Answering this question requires developing a "neutral theory" for language evolution. In this case, the analog of mutation is the variation that is produced as a consequence of failed transmission of languages through the "learning bottleneck" produced by the fact that learners only observe a finite number of utterances. Iterated learning models thus provide a starting point for developing a neutral theory, and understanding which properties of languages can be produced by iterated learning and which properties cannot thus constitutes an interesting direction for future research.

Replicators

This discussion opens up the question of whether or not we should be thinking about cultural transmission/social interaction models in terms of competition among replications or, more excitingly, in terms of different levels or types of replication. It is tempting to propose (as outlined in Kirby 2006) that the emergence of syntax marks a change from one type of replicator (solitary replicators) to another (ensemble replicators), to use terms from Szathmáry (2000). However, this raises the issue of what exactly these replicators are, and whether their dynamics are best described in terms of selection at all.

It appears that the answers to these questions vary enormously depending on one's perspective on the best way to represent the knowledge being acquired/adapted by individuals and the mechanisms for acquiring that knowledge. For example, one view of language might propose that we internalize a set of constructions (e.g., Croft 2000) that have a fairly straightforward relationship with utterances. Accordingly, we might reasonably think of these constructions as replicators, with selection being driven by speakers choosing among constructions to use to produce an utterance. Alternatively, we could think of learners as providing selection pressure, with the constructions that produce the most evidence for their existence in the data available to the learner ending up being the most stable through the learning bottleneck. Here, we can imagine constructions competing for place in the learners' input.

Another view might be that a language is a hypothesis for which we select on the basis of evidence combined with an inductive bias. Where are the replicators here? Who is doing the selection? Given this latter perspective, the neutral model outlined earlier appears more appropriate.

Which of these perspectives is correct? It is possible that in fact they are compatible—that they are different ways of analyzing the same process, namely social/cultural *adaptation*. The challenge is in seeing how these analyses relate to one another and to the models that exist in the literature.

Incidentally, we need to be clear that when we refer to selection and replication here, we are not talking about selection of heritable genetic variation (although that is clearly relevant to language evolution and to models of language evolution). Nor are we referring to the natural selection of cultural variants, a mechanism by which fitter individuals are more likely to survive and pass on their cultural traits (although this, too, is likely to be important). Instead, we are talking about the kind of adaptation that occurs purely through the complex process of a repeated cycle of utterance creation, interpretation, and internalization that happens in language transmission—be it in an iterated learning model focusing on vertical transmission or a negotiation model focusing on social coordination.

Gene–Culture Coevolution

As we pointed out at the onset of this chapter, the current focus of language evolution modeling lies in cultural evolution. It is, however, clear that a complete picture must integrate cultural with biological (genetic) evolution. Formal modeling of gene–culture coevolution began in the mid-1970s (Cavalli-Sforza and Feldman 1981). Briscoe (2003, this volume) reviews models of gene–culture coevolution applied to language evolution. The basic idea is to use the population geneticists' recursion equation formalism for cultural as well as genetic evolution. The result is a system of linked dynamic equations that keep track of genes and culture as they change through time. In general, genes can influence culture via decision-making rules. An innate syntax might constrain

the evolution of languages. The flow of causation will also generally work the other way. An element of a culturally transmitted protolanguage might exert selective pressure on the genes. If genetic variation exists in the innate supports for language, and if more efficient communication is favored, the variants that make the protolanguage more sophisticated will increase in the population. Since cultural evolution will tend to be faster than genetic evolution, cultural evolution will tend to be the driving partner in the coevolutionary circuit and genetic evolution the rate limiting step. This does not tell us anything about the division of labor between genes and culture at evolutionary equilibrium. That will depend upon many contingent costs and benefits of transmitting adaptations genetically versus culturally. Very broadly speaking, the genetic and cultural subsystems are both adaptive systems, and selection may be more or less indifferent as to how a given adaptation is transmitted.

Although a complete coverage of gene–cultural modeling remains as a task for the future, one question which has already been studied by gene–culture coevolutionists is whether and how the evolution of various human adaptations may facilitate or constrain the evolution of language (see Boyd and Richerson 1985, 1996, 2005). Language would seem to require a large measure of cooperation. Otherwise hearers could not trust speakers. The noncooperative case seems to exemplify the situation for most other species. Hence communication systems in other animals are rather limited. Even in humans, people who live in different societies may not be trustworthy sources of information. Hence the evolution of linguistic differences between human groups may be adapted to limit communications from untrustworthy others.

Advanced Recurrent Neural Network Models

There exist a number of RNN architectures designed to model complex linguistic (or visual) processing that are both computationally powerful and partly biologically plausible. These models have not yet been used as a basis for evolution of language studies. Due to their expressivity and the availability of advanced learning algorithms, they appear to be promising carrier formalisms for future evolutionary studies of grammatical processing.

The SHRUTI family of connectionist architectures (e.g., Shastri 1999) represents a long-standing research strand to explain fast-forward inferences in semantic text understanding. These models are very complex, hand-designed networks of semantic and syntactic processing nodes which communicate with each other by biologically motivated neural spike codes, enabling combinatorial binding of different representation nodes across the network.

In machine learning, a recent landmark paper (Hinton and Salakhutdinov 2006) has unleashed a flurry of research in deep belief networks (DBNs) and restricted Boltzmann machines (RBMs). With these models and novel learning algorithms, it has become feasible for the first time to train deep conceptual hierarchical representations from large-volume real-life data in an unsupervised

way. While this field has been preoccupied thus far with visual learning tasks, the step toward speech/text input is imminent.

Another family of hierarchical RNN-based models for learning representations of very complex multiscale data is emerging. These models arise as hierarchical/multiscale extensions of *echo state networks* (Jaeger and Haas 2004) or *liquid state machines* (Maass et al. 2002), the two main exemplars of a new computational paradigm in neuroinformatics referred to as *reservoir computing* (Jaeger et al. 2005). They share a number of important characteristics with the neural models of speech recognition explored by Dominey (2005; Dominey et al. 2006). In this field, language and speech modeling is indeed an important target domain.

An important characteristic of all RNN models, which sets them apart from symbolic grammar formalisms, is that speech/language processing is construed as a fast, self-organizing dynamical system, which does not need search subroutines and does not build interim alternative interpretations. On the positive side, this leads to very fast processing (timescale of a few neuronal delays); on the negative side, if an interpretation trajectory goes astray, this has to be detected and separate repair mechanisms have to be invoked.

Creatures-based Modeling

Current simulation-based studies on language change concentrate primarily on cultural transmission dynamics. Neither brain structures nor genetic determinants for such brain structures are modeled. This makes simulation-based research blind to some of the questions that are raised in biological evolutionary accounts of the origins of language (in this volume, see chapters by Givón, Fedor et al., Hilliard and White). A potentially powerful avenue would be to simulate brain–body coevolution along the lines staked out in Artificial Life and Evolutionary Robotics (e.g., Sims 1994; Nolfi and Floreano 2000; Szathmáry 2007). In this research, artificial creatures are evolved in simulation or in physical robotic hardware. A creature has a body equipped with sensors and actuators, and is controlled by an artificial neural network that coevolves with the body. Research of this kind has resulted in the evolution of surprisingly complex and adaptive behavior repertoires driven by neurocontrollers of surprisingly small size. However, linguistic capabilities have so far largely fallen beyond the scope of this research (cf. Wischmann and Pasemann 2006). It appears a natural and fascinating endeavor, albeit computationally challenging, to implement simulation scenarios where body+brain creatures are evolved under selective pressures that favor efficient communication. In this way, simulation-based research might tell an almost complete (if duly simplified) story, connecting mechanisms of genetic coding of neural structures and the ensuing slow "biological" adaptations with the fast cultural transmission dynamics that are the hallmark of today's investigations.

Detailed Models of Language Learning During Development

Most models of language and syntax evolution treat an individual's learning of language during their "childhood" in a very simplistic fashion. However, the transmission of language from one generation to the next is clearly a central aspect of the evolution of language. Thus, more elaborate modeling of the acquisition of language during infancy and childhood is needed. Ideally, such models would take the embodied nature of language learning into account, capturing how the learner interacts with their physical and social environment. At the same time such models should be constrained by developmental psychology and developmental neuroscience, providing constraints regarding the underlying neural structures, representations, and learning mechanisms, as well as the nature of the language input to which infants are typically exposed. Thus far, such approaches have been mostly restricted to learning early precursors of language, such as gaze following (Triesch et al. 2006) or learning of word meanings (Yu et al. 2005), but the time seems ripe to extend such models toward the acquisition of grammatical structures.

Case Studies

Scientific fields often organize their research around key challenges that are accepted by a large group of researchers independently of the methods they are using. In technical fields, such as machine learning, robotics, or high performance computing, there are often well-defined challenges against which different research groups compete, often leading to very rapid progress (e.g., as seen in the Robocup). What would such key challenges look like in the case of language evolution? One possibility is to pick a certain domain which has been grammaticalized in many languages of the world, although often in different ways, and show what cognitive mechanisms and interaction patterns are needed to see the emergence of such a system in a population of agents. Another possibility is to develop evolutionary models that are also capable of capturing psycholinguistic data on actual human language behavior.

In terms of domain, consider the following example: Many languages have the means to express predicate–argument structure through a system of case grammar, either expressed morphologically or through word-order structure. The emergence of such a system requires not only the emergence of conventions but also the emergence of the semantic and syntactic categories that underlie it. A lot of data is available from historical linguistics showing how such systems have arisen in human natural languages, often by the grammaticalization of verbs, and these data could be used to constrain potential models. Some attempts have already been made to explain the emergence of case grammar, from the viewpoints of each of the paradigms introduced earlier (Moy 2006; Jäger 2007; Van Trijp 2008), and these can act as a starting point for tackling this challenge. It is not difficult to imagine other aspects of grammar that

could form the focus of well-defined challenges, and once these are addressed, more challenging ones could be attempted, such as clause structure with long-distance dependencies.

Conclusion

Mathematical and computational models of language evolution make it possible to examine the consequences of certain theoretical assumptions by mathematical deduction, large-scale computer simulation, or robotic experimentation. Several efforts are under way to apply this methodology to questions related to the problem of the origins and evolution of language. No single paradigm or methodology exists; instead, multiple paradigms explore different questions. At present, models focus primarily on the origins of lexicons, categories that can act as building blocks for conceptualization, and simple languages with few of the complex structuring principles found in human languages (cf. Briscoe, this volume). However, we are confident that the technological foundations and mathematical tools will become progressively more sophisticated and thus be able to tackle increasingly deeper and more intricate questions relating to the origins and evolution of syntax in language.

Glossary

ADV	Adverb
AF	Articulate fasciculus
AFP	Anterior forebrain pathway
AG	Angular gyrus
AMH	Anatomically modern humans
AP	Adjective phrase
ApoE4	Apolipoprotein E
BA	Brodmann's area
Baldwin effect	Proposed by James Mark Baldwin, the Baldwin effect, also known as Baldwinian evolution or ontogenic evolution, consists of a mechanism for specific selection for general learning ability. Selected offspring tend to have an increased capacity for learning new skills rather than being confined to genetically coded, relatively fixed abilities. Suppose a species is threatened by a new predator and there is a behavior that makes it more difficult for the predator to kill individuals of the species. Individuals who learn the behavior more quickly will obviously be at an advantage. As time goes on, the ability to learn the behavior will improve (by genetic selection), and at some point it will seem to be an instinct, thus the repertoire of genetically encoded abilities can ultimately increase. This ultimate phase of the Baldwin effect is related to the concept of genetic assimilation.
BG	Basal ganglia
Binding	Relationship between anaphors and their antecedents
Birdsong	Bird vocalization includes both bird calls and bird songs. The distinction between songs and calls is based upon inflection, length, and context. Songs are longer and more complex and are associated with courtship and mating, whereas calls tend to serve such functions as alarms or keeping members of a flock in contact.
Birdsong syntax	Temporal sequence in which discrete units of song (notes, syllables, phrases, motifs, bouts) are produced
BOLD signal	Blood-oxygen-level-dependent signal
CFG	Context-free grammar
CHiP	Chromatin immunoprecipitation
CNTNAP2	Contactin-associated protein-like 2 gene