Self-organization and Selection in Cultural Language Evolution

Luc Steels

This paper is the author's draft and has now been officially published as:

Steels, Luc (2012). Introduction. Self-organization and selection in cultural language evolution. In Luc Steels (Ed.), *Experiments in Cultural Language Evolution*, 1 – 37. Amsterdam: John Benjamins.

Abstract

This chapter outlines the main challenges a theory for the cultural evolution of language should address and proposes a particular theory which is worked out and explored in greater detail in the remaining chapters of this book. The theory rests on two biologically inspired mechanisms, namely selection and self-organization, mapped onto the cultural, more specifically, linguistic domain. Selectionism is an alternative to rational top-down design. It introduces a distinction between processes that generate possible linguistic variants in a population (for example, different ways to express tense and aspect) and processes that select some variants to survive and become dominant in a language, based on criteria that translate into increased communicative success, such as expressive adequacy, minimal cognitive effort, learnability and social conformity. Self-organization occurs when speakers and hearers align their communication systems based on the outcome of each interaction. It explains how convergence may arise without central coordination or direct telepathic meaning transfer. This chapter explains these basic hypotheses in more detail and introduces a methodology for exploring them based on the notion of a language game.

1. Introduction

A general theory of language evolution should cover the origins and evolution of language from three perspectives: biological, social and cultural.

1. *Biological*: Language is learned, produced and comprehended by human brains to handle communicative challenges arising from the ecological conditions in which humans live. A first important question is therefore how the biological structures and processes needed for language could have evolved in the course of biological evolution. This question is currently being studied using the same Darwinian framework as employed in other areas of evolutionary biology, namely, by investigating which cognitive functions need to be present for language, which neurobiological structures implement them, what the genetic bases of these structures are, whether there are any homologs or precursors in non-human primates or even other species, and when and how the genetic bases for language could have evolved in human evolution. Recent survey papers exploring this biological point of view can be found in (Bickerton & Szathmary, 2008).

2. *Social*: But language is also a social phenomenon. It appears only in the context of certain cooperative interaction patterns, social relations, and tasks for which the use of symbolic communication may be beneficial. A second important question is therefore to clarify the nature and origins of the social bases of language, for example, by identifying the kind of ecological changes that could have triggered the need for open-ended forms of sophisticated cooperation or by explaining how the level of trust necesssary for joint attention and symbolic communication could have become accepted in early homonin societies (Knight et al., 2000). The tools of anthropology and social science are more appropriate for tackling this second question. It may involve, for example, a study of hunter-gatherer societies to understand the possible earliest role of symbolic communication (Knight et al., 2001), or a study of the limitations of social interaction in non-human primates (Tomasello, 2003).

3. *Cultural*: Nobody doubts that languages change. It suffices to compare an English text from a few centuries ago with contemporary English (Aitchinson, 2001) or to compare a creole with its source language (Thomason & Kaufman, 1988), and the reader immediately observes that no aspect of language is immune to change, whether it is at the level of speech sounds, sound patterns, morphemes, words, syntactic patterns, semantic structures, meanings or pragmatic turn-taking patterns. The changes concern both the appearance of new linguistic elements and the disappearance of existing ones, sometimes remaining only as fossilized recycled remnants. So language is a cultural phenomenon as well, that is undergoing steady evolution at a cultural level. A third important question for a general theory of language evolution is therefore how this cultural evolution has been operating. Moreover, when we apply the uniformitarianism principle, namely that the cultural processes cur-

rently observed in language change have been operating since the origins of the first language, then we can in principle roll back history and understand the origins of the first forms of language as well (Heine & Kuteva, 2008).

Clearly these three perspectives strongly interact, as one perspective introduces constraints on what assumptions make sense from another perspective. For example, theories of cultural evolution cannot rely on social conditions which are unrealistic (such as centralized control over language) or on cognitive capacities which humans do not have (such as direct telepathic meaning transfer). Moreover, increased complexity in one area pushes complexity up in another, resulting in a spiraling process (Figure 1): Increased social and ecological complexity requires increased language complexity which in turn requires increased brain capacities which then further support increased social complexity. Going in the other direction: increased language complexity supports increased social complexity which puts increased pressure on the development of mental processes that then in turn again support greater linguistic complexity.

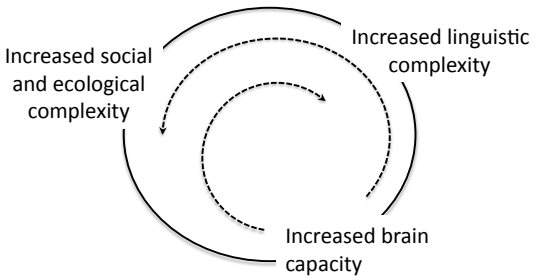

Figure 1. *The different forces shaping the emergence of language push complexity up in a spiraling process.*

Despite of these strong interactions, we focus in this book on a theory for the *cultural* evolution of language only. Such a theory should not only allow us to explain in general terms why languages change, or why and how complex vocabularies and grammatical structures may appear, but it should also allow us to explain specific phenomena of human languages. For example, how and why a case system with the same sort of sophistication as found in Spanish might arise, how a tense-aspect system like we see in Russian could have emerged, how different parts of speech (as for the English noun, adjective, article, adverb) could develop and play a role in constituent structure, how a set of posture verbs could evolve and then be reused metaphorically for talking about positions and movements of objects (as

seen in Dutch), how agreement systems to signal the grouping of the elements in a nominal phrase could emerge (as in Polish), and so on. It is only when a substantial number of concrete cases has been satisfactory dealt with that we can begin to get any confidence in the general principles proposed by a candidate theory.

The goal of this chapter is to inquire first about some of the facts known about human language change and the challenges that any theory of cultural language evolution should address given these facts. It then proposes a particular theory based on linguistic selection and self-organization. The final section introduces an experimental framework based on multi-agent simulations of language games with which one can test this or other theories of cultural language evolution. This framework has by now been used in a large number of case studies, some of them reported in subsequent chapters of this book.

2. Facts about Cultural Language Evolution

2.1. Language Systems and Language Strategies

When studying the historical evolution as it occurred in human languages, we can see change at two different levels: that of language systems and that of language strategies. *Language systems* (sometimes called paradigms) capture the systematicity observed in some domain of the vocabulary or grammar of a language, for example, a system of basic color terms, tense-aspect distinctions, movement verbs or articles. Language systems group a set of paradigmatic choices both on the side of meaning (the conceptual system) and on the side of form (the linguistic system). The *conceptual system*, sometimes called the ontology, includes pragmatic and semantic distinctions that are expressable in this language system and can therefore be used as building blocks for conceptualization. The *linguistic system* includes all the syntactic, morphological and phonological categories and structures to turn a conceptualization into a concrete utterance.

A given language comprises thousands of language systems, which are tightly integrated. Here are some examples:

1. German features a case marking system based on four paradigmatic cases: nominative, accusative, dative, and genitive, which semantically relate to the role of participants in the event introduced by the verb (such as agent, beneficiary, location and instrument). German requires agreement for case, number and gender between nominals and determiners and marks these features morphologically. This is illustrated in the sentence ”Der Hund beisst den Mann”,

where ”der” and ”Hund” is nominative, masculine, singular, and ”den” and ”Mann” are accusative, masculine, singular.

2. Spanish features a system of basic color terms including ”blanco” (white), ”negro” (black), ”rojo” (red), ”verde” (green), ”amarillo” (yellow) , ”azul” (blue), and ”marron” (brown). These colors carve out distinctive regions in the three-dimensional color space formed by the two color opponent channel dimensions (yellow-blue and red-green) and the lightness dimension (dark-light).

3. Russian features a refined aspect system based on ‘Aktionsarten’. Semantically, Aktionsarten highlight segments of the internal structure of the action or event denoted by the verb. For example, *delimitative* emphasizes the beginning and the end, *ingressive* the beginning, *terminative* a definite ending to the action. Aktionsarten are expressed as morphological markers attached to the verb.

Many different language systems are called upon in the production of a single sentence. Consider the example:
“Mario saw a big elephant appearing left of his car.”
This sentence uses a system of lexicalized event types (“see”, “appear”) and object classes (“elephant”, “car”), a system of argument structure to introduce the roles of the two participants (the subject and the direct object), a system of proper names (“Mario”) to refer to a specific individual as well as a system of pronominal reference to refer back to individuals introduced earlier in discourse (“his”), a system of spatial relations (“left of”) and sizes (“big”) and a tense-aspect system to convey that the seeing event happened completely in the past (“saw”) and to highlight the beginning of the appear-event (“appearing”).

The distinction between different language systems does not imply that vocabulary or grammar is stored in terms of separate language systems. This is unlikely. It is now widely accepted that language users store knowledge about the vocabulary and grammar of their language in terms of *constructions*, which associate aspects of meaning with aspects of form (Goldberg, 2003; Michaelis & Lambrecht, 1996; Kay & Fillmore, 1996; Croft, 2001; Steels, 2012). A single construction typically packages a lot of constraints together for efficient parsing or production, and this implies that a single construction may incorporate aspects of several different language systems. For example, a determiner-nominal-phrase construction, as illustrated with the phrase “the mouse”, not only concerns hierarchical structure and word order,

signaling how the meanings of the constituents are to be linked to form the meaning of the whole, but also aspects of determination (definite vs. indefinite), number (singular vs. plural) and agreement (in number and countability between the article and the noun). Moreover, the implementation of a particular language system may have an impact on many different constructions. For example, if an agreement for syntactic features such as gender is used, then lexical constructions have to introduce this information for each word, and features have to percolate in hierarchical structure, for example from the noun to the nominal phrase to handle agreement with the verb or predicate of the sentence.

Typologists call the approach underlying a language system a *language strategy*. For example, they talk about color term strategies, relative-clause formation strategies, case strategies, coordination strategies (for combining nominals), negation strategies (for expressing negation), and so on. Knowledge about a language strategy requires both a meaning component for handling the formation, learning and alignment of the relevant conceptual system and a linguistic component for handling the formation, learning and alignment of the related linguistic system.

There is clear variation across languages in terms of which language systems, and hence which language strategies, they use. For example, Japanese does not have a case system; English does not have a grammatical Aktionsart system; Russian does not have a system of articles; French does not have a system of classifiers (as found in Bantu languages). And even if two languages employ the same strategy, the details of the language systems built with this strategy may still significantly differ. For example, Polish uses a genitive in situations where German uses a dative or accusative; Russian has an instrumental case absent from Spanish; Hungarian features cases expressing place or direction, such as inessive (in/inside), adessive (at), or illative (movement towards the inside). Indo-European languages typically express the same relations with prepositions.

By mapping out different languages, strategy trees can be created, such as the one shown in Figure 2, where different possible strategies are identified for each major domain in language followed by different instantiations for each of these strategies. Languages typically combine different strategies for the same domain. For example, although English now uses constituent order and prepositions as the main vehicle for expressing the role of participants in events, there are still some remnants of an older case system which shows up in the declension of pronouns.

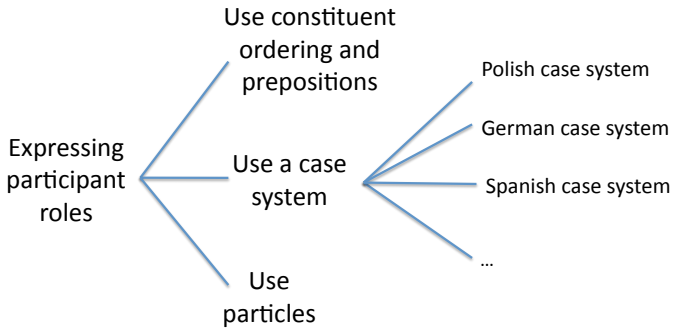

Figure 2. *There are many different ways to express the roles of participants in events: using a case system (as for example in Polish), using constituent order and prepositions (as in English), using a particle system (as in Japanese), and so on, and for each of these strategies many different instantiated language systems can be found in the languages around the world.*

2.2. Language Change and Language Evolution

We know from the historical record that there is ongoing change in the paradigmatic choices of language systems, both at the conceptual and the linguistic level. The change takes the following forms:

1. There can be an increase in complexity of which choices have been adopted. For example, the basic color term ”orange” entered the English language only in the beginning of the 16th century. Before that period, the same hue was referred to as yellow-red. But there can also be a decrease in complexity. For example, Dutch has lost the distinction between dative and accusative case. More generally, there has been a tendency in Germanic languages to simplify existing case systems.

2. The semantic territory that a particular paradigmatic choice covers can shrink or grow at the expense of other choices. For example, when the word ”orange” came into vogue, it pushed aside the regions in the color space covered by red and yellow. Currently, the German dative is overtaking many of the uses that the genitive still has, and the genitive may disappear from German except with possessives, as it has in Dutch.

3. The way a semantic or grammatical feature gets marked can change. Very often one word may evolve into a morpheme attached to another one. For example, the ending "-ed" for English past tense originally arose from the verb "did" following the verb stem, as in "he walk-did", which then became more compact gradually turning into a suffix, as in "he walked". A morpheme may further erode until it gets too weak or even disappears and needs to be replaced by another expression of the same information. This is for example how the negation particle "pas" in French became obligatory (as in "je ne veux pas" (I do not want)), after the Latin source "non" had eroded to "ne". "ne" is occasionally left out entirely in contemporary spoken French and its reinforcer "pas" (originally coming from "un pas" (a step)) has become firmly established as the negation particle.

All these changes are well known and have been reported extensively in the diachronic literature (Heine & Kuteva, 2002). They are significant but do not destroy the logic underlying a particular language system, in other words, the underlying language strategy stays the same.

In addition, we see in the historical record, particularly in creole formation (Mufwene, 2001), that there are occasionally more profound innovations, in the sense that new language strategies may become adopted by the population and others may disappear. Here are a few examples:

1. Words like "yellow" or "blue" used to be brightness terms in old English and have now become hue terms. Speakers of old English predominantly used a color system based on brightness distinctions (now expressed with words like bright, shiny, dull, etc.) which was overtaken at the time of Middle English by predominant use of a system based on hue (with yellow, blue, red, green, etc.). A similar strategy shift took place in many languages (MacLaury, 1992).

2. In Old and Middle English, negation was standardly formed by adding the particle "not" after the main verb, as in "I come not". From the 15th century a variant arose to insert the auxiliary "do" before the verb and add the particle "not" after the auxiliary, as in "I do not come". The two variants existed side by side until the construction with "do" became the norm at the beginning of the 19th century.

3. Latin did not feature a system of articles for determination (i.e. for expressing how the referent of a nominal phrase must be accessed given a class of objects delineated by the nominal) but all languages derived from Latin (French,

Catalan, etc.) developed such a system, usually out of demonstratives (Diessel, 1999).

4. English used to have a case system for expressing the role of participants in events which was comparable in complexity to the case system of Latin or Greek, but this system eroded by the time of Middle English and got replaced by a system primarily based on constituent ordering and prepositions. (Van Kemenade, 1987).

Thousands of such paradigmatic competitions are going on in languages at any given point in time, and often the fact that one variant becomes dominant may impact choices for other variants.

There is occasional debate about whether we should speak in the case of language about change only (as in geology) or whether it is appropriate to talk about true evolution (as in biology). Clearly, when a new strategy originates in a language community and gives rise to a newly emergent language system, it is appropriate to talk about language evolution. The level of novelty and innovation is certainly comparable to true innovation in biology. Even changes in a language system can be very significant and may lead to a ripple effect destabilizing other language systems and eventually requiring the introduction of a new strategy in the language.

2.3. Semiotic Dynamics

The data about language change collected by historical linguists and sociolinguists have also brought up some important facts about how variants, both for language strategies and for paradigmatic choices of a language system, propagate in a population. Language change clearly does not happen instantaneously but is gradual. There is typically a period where novel strategies, meanings, and forms appear as small-scale micro trends, before they start to propagate very rapidly as they are being picked up by the majority of the population. This phase is usually followed by a slowed down propagation rate. Data collected about the frequency by which a population adopts a novel form therefore shows a so called S-shaped curve (Bailey, 1973). This curve is very familiar from population biology, for example in the spreading of genes or diseases. It is also observed in the spreading of innovative products or new forms of government, suggesting that the propagation of novel language meanings or forms is based on the same universal laws as found in all complex adaptive systems (Holland, 1992), raising the hypothesis that language should be viewed as a complex adaptive system itself (Steels, 2000).

A new norm can itself be overthrown again when yet another new linguistic innovation appears, similar to the way that new technological inventions may disrupt existing products, even if they have complete dominance of a market (Christensen, 1997). Moreover, older forms may still stick around, as we see for example in English past tense formation which shifted from a once dominant variant in which the vowel of the verb stem changed (as in "spring/sprang") to the currently dominating solution using the morpheme "-ed" (as in "walk/walked"). Often a cooperation between different strategies emerges. For example, when a hue-based color system became dominant, it did not entirely annihilate the existing brightness-based system (even though it was able to "steal" a lot of brightness words to become hue words), rather, the two strategies now co-exist side by side and can be combined, as in "shiny yellow" which makes use of both brightness and hue terms.

Interestingly, speakers are usually aware that several competing paradigmatic language system choices or language strategies are circulating in their language community and they know which social, age, or geographic characteristics are associated with them. They are able to comprehend utterances based on competing solutions even though they themselves prefer another one, conform to the geographic or social group they belong to. Moreover, it is not necessarily that the population converges entirely towards the same solutions and individuals may in certain contexts still stick to older forms. So multiple outcomes can remain side by side.

A good example of this kind of semiotic dynamics is the current evolution within the Spanish pronoun case system (see the later chapter (van Trijp, 2012)). Accusative and dative cases are collapsing and persistent variation is found for three competing variants for expressing the surviving case: *leismo* with *le* instead of *lo*, *laísmo* with *la* instead of *le* when there is a feminine referent, and loísmo with *lo* instead of *le* for masculine or neuter referents. Each of these competing forms is surviving and speakers are familiar with the different variants and can even imitate them. The variants are associated with specific regions in Spain and speakers can identify from which region a speaker comes based on which variant he or she is using.

2.4. Challenges for Theories of Cultural Language Evolution

Given these observations, we can now define more clearly the fundamental questions that a theory of cultural language evolution should address.

First of all, it should explain *how a language system may emerge and continue to change*, assuming that all individuals of the language community share the same

strategy. For example, the theory should be able to explain how a basic color term system, or a case system, or an agreement system can arise and continue to evolve, supposing that all language users know and use the same relevant language strategy. The paradigmatic choices both on the conceptual and on the linguistic side are open to change but the basic systemic principles are fixed and shared across all individuals in the community.

The assumption that everybody in a language community shares *a priori* the same language strategies is of course unwarranted, given all the cross-linguistic variation and change that we see at the level of language strategies as well, and so a theory of language evolution should also attack a second, more difficult goal, namely to explain *how a new language strategy can emerge and propagate in a population*, how it can enter and possibly win a competition against other strategies, or on the contrary develop a symbiotic relationship.

What form should answers to both questions take? On the one hand, we will need to posit some general principles on how structure may arise in language. I will argue soon that linguistic selection and self-organization can play this kind of role. But then we need to instantiate these principles by identifying what kind of cognitive functions and interaction patterns are needed to put them into action. The cognitive functions have to be specified at a sufficiently concrete level so that the theory's adequacy for explaining empirically observed linguistic phenomena can be objectively tested, just like the theory of evolution by natural selection needs to be instantiated with concrete facts about physics, chemistry, ecology, genetics and development in order to explain a particular trait such as lungs or butterfly wings. Note that these explanations go beyond a description of the historical facts on how a particular linguistic meaning or form has arisen in a language community, because they try to *explain* why these facts have happened.

Third, a theory of language evolution should explain the *semiotic dynamics* we see in cultural language evolution, such as the S-shaped curve. By semiotic dynamics, we mean the evolution over time of various macroscopic properties of a language or its use in a community, such as, how widespread a particular linguistic convention has been adopted, the average size of the vocabulary of all individuals, the average number of distinctive basic color concepts, the similarity in grammatical constructions for expressing argument structure, how dominant a particular strategy is.

Semiotic dynamics explanations should follow from proposed answers to the previous questions, in other words from the cognitive functions proposed as necessary and sufficient for inventing and coordinating language systems and language

strategies. It is certainly possible to model the S-shaped curve with an equation (namely the logistic or Vanderhulst equation) using aggregate quantities, such as the frequency of a particular form in the population at a particular point in time, but such models do not have any explanatory force from the viewpoint of a theory of language evolution. They merely describe the mathematical structure of the phenomena we observe. Just as the use of a logistic equation for describing the propagation of a disease does not explain what the disease is, how it is caused, how it originated, how it transits from one individual to another, or what counteraction could be taken.

3. Language Evolution through Linguistic Selection

There has been considerable research during the past few decades to develop theories of cultural evolution, see for example, Cavalli-Sforza & Feldman (1981); Dawkins (1976); Boyd & Richerson (2005); Jablonka & Lamb (2005), and these theories have formed the background for several research programs in language evolution, such as the lingueme theory (Croft, 2000). Nevertheless, there is certainly no consensus yet about which candidate theory has the most explanatory force and many open questions remain.

This section outlines the *selectionist theory of language evolution* that will be elaborated and tested extensively in the rest of this book. The theory introduces selection at the level of language systems and language strategies (as in Mufwene (2001); Heine (1997)) and uses self-organization based on alignment to explain how a population may converge to sufficiently shared systems and strategies to make communication possible (as in Garrod & Doherty (1994); Garrod et al. (2010)). Because it relies on the ability of language users to recruit cognitive functions and conceptual frameworks in the service of language, it is also known as the *recruitment theory* of language origins (Steels, 2007).

Just consider, as a thought experiment, that a group of intelligent agents is faced with the task of coming up with a communication system. The task is formidable for a wide variety of reasons. Even if we assume that speech is the medium, there are still a virtually infinite set of possible ways in which a language could be constructed. Moreover, it is not just a matter of evolving possible forms (words, morphemes, grammatical constructions). The language users must also come up with the conceptualizations of reality that they are going to use for language, and these meanings must be grounded in reality through their sensori-motor systems.

Why is all this so difficult?

1. Human languages must be capable to talk about anything. As it is impossible to know in advance all the different issues that may come up, languages must be open-ended. And this implies that language users must be allowed at all times to expand their language for dealing with new meanings or to introduce new forms to cover these meanings.

2. There is no central coordinator nor powerful committee that designs human languages in a top-down fashion and imposes them by decree. Everyone has equal rights to adapt the language to their aims. Massive variation is therefore unavoidable, particularly because one individual cannot know what others have already invented.

3. Language users cannot inspect or change the minds of others directly in a kind of telepathic way. So hearers need to make guesses about the meanings of forms they do not know yet and these guesses could be wrong. Speakers may make the wrong assumptions about what hearers know, and this cannot only lead to misunderstandings or failures in communication but also to growing divergence of their language systems.

These features make it impossible to design human languages in a top-down rational way, even though this has been tried several times throughout human history (Eco, 2000). An alternative to top-down rational design is selectionism. Selectionism splits the task of coming up with a design in two: There is a process of *generating* possible variants, often by making small changes to existing variants, and there is a separate process of *testing* whether these variants satisfy desired selection criteria. Crucially, there is also a *self-enforcing causal loop* between the outcome of selection and the frequency with which a variant is maintained and used to generate new variants. This makes a selectionist process cumulative: partly working solutions can be built upon further to come up with even better ones.

Selectionism was first proposed by Darwin and his contemporaries in order to explain the origins of species (Maynard Smith & Száthmary, 1993). A rational top-down design of organisms indeed faces the same difficulties as the design of languages: New challenges come up due to changing environmental and ecological conditions, such as a rise in temperature or the arrival of a new predator. This requires a species to adapt constantly to novel conditions which could not be foreseen 'at design time'. Biological evolution solves this problem by generating variants through mutation or recombination of DNA or by taking aspects of the environment into account during development, and by testing these variants in terms of their

survival within the ecosystem. The self-enforcing causal loop goes through reproductive success. If a variant leads to more offspring, its frequency increases in future populations and it can therefore form the basis for possible further improvements.

Today, selectionism is recognized as a general principle that can and has been applied to many different types of systems, such as in chemistry (Luisi & Oberholzer, 2001), neurobiology (Edelman, 1987) or economics (Arthur, 1996). It has also been used successfully to come up with new artificial systems using the paradigm of genetic algorithms (Koza, 1992), for example to evolve controllers for robots (Nolfi & Floreano, 2000). It is therefore not so surprising that selectionism can also be mapped to the cultural and more specifically to the linguistic domain, in which case the nature of the possible variants, their generation and maintainance in a pool of variants, the selectionist criteria, their application, and the self-enforcing causal loop all have to be transposed to the linguistic level.

3.1. Linguistic Selection Criteria

Let us first reflect on what kind of selectionist forces are relevant in the case of language. Clearly, the bottom line is that language users want to achieve **communicative success**. Communicative success means that the speaker achieves his or her non-linguistic goal. For example, if a speaker wants a hearer to perform some action (such as "Can you get me the car keys?"), there is communicative success if the hearer indeed performs the requested action. Speakers and hearers generally do not care whether their sentences are perfectly grammatical or complete or semantically accurate. Speakers assume that hearers are intelligent enough to infer the missing information and that they can fill in the required details from the context (Sperber & Wilson, 1986). In normal spoken dialogue, communicative success is immediately obvious, either from the actions that the hearer performs as a consequence of an utterance or by gestural cues whereby the hearer informs the speaker that he or she is being understood.

Communicative success rests on a number of features of a language which thus act as the key selection criteria:

Expressive adequacy: The language systems available to the speaker and the hearer must have the necessary expressive adequacy to reach communicative success. The available conceptual systems must include the needed conceptual distinctions and the available linguistic systems must be able to express these distinctions.

Cognitive effort: In order to cope with the incredible speed of normal language production and comprehension, reducing cognitive effort is of primordial impor-

tance. Cognitive effort is expended at all levels of language: How much time and memory needs to be spent in coming up with adequate conceptualizations? How complex is process of constructing a sentence? How difficult is it to articulate the speech sounds? How efficiently can the sentence be parsed? How complicated is the interpretation of the sentence? If sentences are too complex to comprehend, for example because they trigger a large amount of combinatorial search, then hearers give up. Or if sentences are too hard to produce, the production process will be too slow to maintain the hearer's attention.

Learnability: Speakers must occasionally expand their language systems to express novel meanings or better capture the attention of hearers with a 'fresh' way of saying something. But these innovations will lead nowhere if they cannot be acquired by hearers. They will lead to communicative failure and the innovation will not propagate in the population and therefore it has no chance to survive in the shared language.

Social conformity: Speakers and hearers can greatly optimize their language production and comprehension by making their language systems as similar as possible. Moreover, details of language use, such as pronunciation, signal to which group a speaker belongs, and so speakers seek social conformity with the language use of their peers in order to belong or they try to behave in a similar way as the most prestigious group. The opposite happens as well. Speakers try to mark themselves as linguistically different from a group to which they do not want to belong.

Language users certainly cannot consciously construct their language to satisfy these selectionist criteria for the reasons given earlier, even if they would want to. They can have no foresight about all the possible distinctions that are going to be relevant in their world because new objects, new artefacts, new types of interactions always come up. They have no obvious way to know consciously how much cognitive effort a particular sentence is going to require from the hearer. They cannot know in advance whether a metaphorical extension will be understood or whether the coercion of a word into a new grammatical function is going to be grasped by the hearer. And they cannot know for sure what the norm is in the group because they have no general overview, but can gain evidence only from local interactions with others whom they might not even be sure of whether they belong to the same group. The only thing they know for sure (and not even always or immediately) is whether a communicative interaction was successful. This is why language evolution has to be a selectionist process. Language users might be able to generate or reuse certain variants but they can never be absolutely sure that these variants satisfy the linguistic selection criteria leading to persistent communicative success.

3.2. Linguistic Selection of Systems and Strategies

We have seen that empirical observations of historical language change point to a change at two levels: that of language systems and that of language strategies, and hence linguistic selection is expected to operate at each level as well. In the case of language systems,

1. Language users store the paradigmatic language system choices that they introduced or acquired using a particular language strategy in their own memory and choose what they believe to be the best solution to deal with a particular communicative situation.

2. The various selectionist criteria discussed earlier (adequate expressive power, minimal cognitive effort, learnability, and conformity to the social group) intervene implicitly in the production and comprehension of a sentence or have an impact on the communicative outcome. For example, if the communication succeeded, the language had enough expressive power to deal at least with the current situation.

3. The self-enforcing causal loop between linguistic selection and the frequency and preservation of a certain language system variant goes through communicative success (Figure 3). If a particular variant was chosen by the speaker and lead to a successful communication, then the participants in this interaction should change their memories in such a way that they are more likely to use this variant in their future interactions. On the other hand, if a failure arose, the future likelihood of using this variant should diminish. If all language users behave this way, we will see convergence at the population level to those variants that achieve the highest communicative success.

As in the case of biological evolution, there is a cumulative effect because solutions that have been successful are maintained in the population, enabling their further use as building blocks for tackling more challenging communicative goals in more demanding contexts.

But where do language strategies come from? We can apply the same selectionist logic:

1. Language strategies are stored in the memories of speakers and hearers and exercised when a language system needs to be expanded or coordinated. Individuals generate variants based on a *recruitment* process. They configure cognitive functions to implement a new strategy from scratch or they start

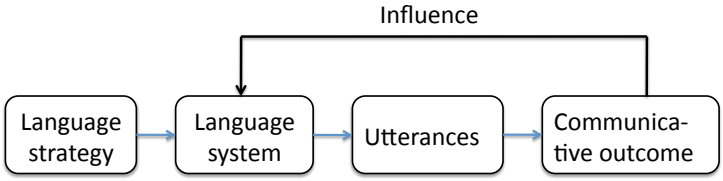

Figure 3. *A language strategy can give rise to a large set of possible paradigmatic choices. Which specific choices are adopted by a language community depends on the outcome of communication. A successful choice propagates and is used more so that it has even more success. So we get a self-enforcing causal loop between communicative outcome and paradigmatic language system choices.*

from an existing strategy and perform operations on it to yield a new variant (exaptation). For example, they may have a strategy for spatial language which they now transpose to the time domain.

2. The selection criteria for a language strategy are based on whether the language system it has helped to build and maintain leads to long-term communicative success, which implies indirectly that the basic linguistic selection criteria (adequate expressive power, minimal cognitive effort, learnability, and conformity to the social group) are optimized.

3. There is a self-enforcing causal loop between linguistic selection and the use, frequency and preservation of a certain language strategy. This loop goes through the long term communicative success of the language system built by the strategy (Figure 4).

Although human language users cannot be sure whether the strategies they use and the system choices they make are going to lead to communicative success, they are not entirely blind either. They possess a powerful cognitive system that can be extremely helpful to make good guesses. This is the crucial difference between genetic evolution by natural selection and language evolution by linguistic selection and explains why language evolution goes so rapidly.

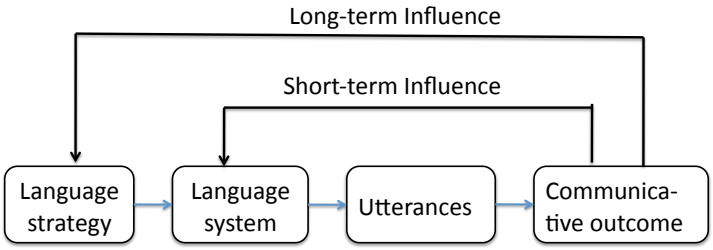

Figure 4. *Many different language strategies are possible. Which strategy is adopted by a language community depends on the success of the language systems the strategy is able to set up and maintain. And the more a particular language strategy was implicated in communicative success, the more frequently it should be used and the more likely it propagates and thus becomes used even more in the population. So we get a second self-enforcing causal loop, between long term communicative success and language strategy choices.*

4. The Role of Cognition

Biological evolution gets physics and chemistry 'for free'. Linguistic evolution also gets something for free, namely cognition. First of all, language is embedded in a larger cognitive system that it can make full use of. For example, spatial categories that are important for navigation can be recruited as possible meanings in verbal communication, the ability to detect sequences or hierarchical structures, which is needed to make sense of the real world, can be recruited to exploit word order or handle constituent structure, fine-grained motor control can be recruited to drive speech articulation. Because these cognitive capabilities are necessary to deal with the world and normal social interaction there is a high chance that they are shared, at least to some extent, with other language users. But cognition can also play an important role for speeding up language evolution by monitoring language processing and its outcome.

To implement this second role for cognition we need a cognitive architecture with two different layers (Figure 5). There is a first layer of routine processing based on the language systems that the speaker or the hearer have available in memory. When there is a straightforward, almost ready-made solution to a communicative challenge, this routine layer is employed and may be entirely adequate. But there

must also be a second layer of *meta-level processes* which monitor what is going on and can possibly take action if a problem arises.

Monitoring is based on *diagnostic procedures* that inspect the structures being built and the outcome of communication and they may flag dead ends or unusual conditions in processing. For example, processing may get stuck when the speaker cannot find the right word to complete a sentence according to the original plan; a word may not be recognized by the hearer due to sloppy pronunciation of the speaker; too much ambiguity may have lead the hearer to explore a path in the search space that turns out to be a dead end; the hearer may be confronted with an ungrammatical sentence fragment so that constituents cannot be put together properly, and so on.

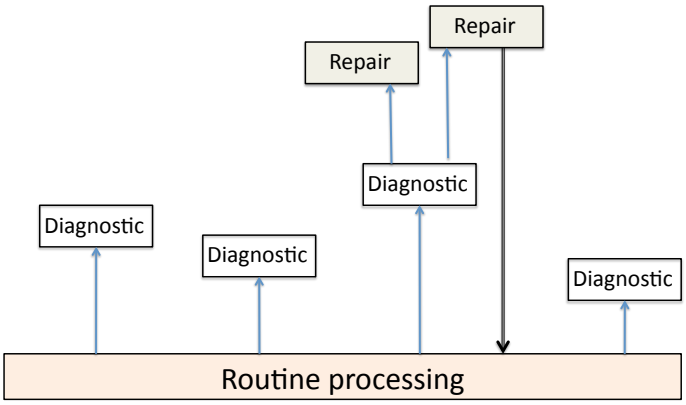

Figure 5. *Routine processing during parsing and production is augmented with meta-level processes that perform diagnosis and in case of problems look for a repair that could handle them. After the repair has been enacted, routine processing can often continue.*

Repairs are actions that speakers and hearers take in order to handle problems signaled by a diagnostic. For example:

- When a word fitting the syntactic context cannot be found in language processing, the speaker can backtrack and come up with a new global sentence plan so that this word is no longer needed.
- When a word is not understood by the hearer, top-down prediction based on meaning inferred from the context can help to fill in this gap in the speech stream.

- When a dead end is reached in the hearer's search space, the hearer can backtrack and explore another path in the search space that seemed at first less likely.

- When the speaker has produced an ungrammatical language fragment, the hearer could allow more flexibility in applying grammatical constructions and still be able to come up with a partial interpretation.

Such repairs are clearly observable in normal dialog (Levelt, 1989) and explain why language processing is so robust (Steels & van Trijp, 2012).

Even more interesting, and more relevant for language evolution, are situations where the current language system is fundamentally incapable to deal in a routine way with the current communicative challenge. In that case, diagnostics and repairs can help to build up or improve the language systems of speaker or hearer according to specific language strategies. Indeed, we can now see that language strategies consist, operationally speaking, of a set of diagnostics and repairs that expand a language system while maintaining the existing systematicity and while trying to push the language towards greater communicative success based on more adequate expressive power, less cognitive effort, increased learnability, and greater social conformity. Here are some examples:

- *Expressive adequacy*: The speaker could be missing a particular spatial distinction in order to distinguish two objects based on their spatial location. A repair might then consist of introducing a new spatial distinction and introducing a new term for it (Spranger, 2012).

- *Cognitive effort*: The speaker could make it easier on himself by abbreviating certain words or by leaving out certain grammatical function words. This would at first lead to more effort for the hearer until he has integrated this new convention into his own language system. The end result will be more compact sentences and faster processing.

- *Learnability*: When the speaker introduces a new word in order to express a particular concept, he could decide to re-use an existing word whose meaning already suggests or includes the new word and rely on the ability of hearers to infer by analogy how the existing word has been expanded (Wellens et al., 2008). Or in order to come up with a word that fits a particular syntactic context, the hearer could re-use an existing word which covers more or less the same meaning, and coerce this word into a new syntactic function. This is

easier to learn than acquiring the meaning and function of a brand new word (Steels & van Trijp, 2012).

- *Social conformity*: The hearer can be more conform the language used by the speaker by storing the solutions that were used by the speaker, even if he does not (yet) prefer these solutions himself. Language systems choices can be marked in memory for the social or regional group with which they are associated and the speaker can then decide on the use of a variant taking this information into account.

Some repairs intended to improve one selectionist criterion may actually run counter to another criterion. For example, leaving out the morphological endings that express case and gender reduces the effort for the speaker but not for the hearer. It may eventually lead to a collapse of the case system, hence a loss of expressive power, which then might trigger the introduction of a new language strategy to express the same information in another way. This phenomenon is believed to have happened when the English case system collapsed and was replaced by a strategy based on constituent ordering and prepositions (Van Kemenade, 1987).

The power of diagnostics and repairs can be greatly improved by *re-entrance* (Steels, 2003). While producing a sentence, the speaker can re-enter the partial or complete sentence back into his own parsing and interpretation system thus simulating the hearer, and while comprehending a sentence, the hearer can re-enter the reconstructed meaning back into his own language production system thus simulating the speaker. Re-entrance is not only useful for repairing ongoing dialog, for example for enhanced flexibility and robustness, but it can also be used for a more intelligent expansion or improvement of the available language systems, particularly to minimize cognitive effort or improve learnability. Here are again some examples:

- *Expressive adequacy*: The speaker may try to parse and interpret a sentence that he is producing in order to detect whether the hearer would be forced to inspect the world model in order to disambiguate the derived meaning. The speaker may then repair this situation by increasing the expressive power of his language, for example by being more precise about the role of participants in the event (van Trijp, 2012).

- *Cognitive Effort*: The speaker may re-enter a sentence he produced into his own parsing system in order to detect whether the sentence might lead to

unnecessary combinatorial search. If that is the case he could introduce additional syntax to dampen the search so that the hearer has a higher chance to know which is the best path to explore (Steels & Wellens, 2006; Beuls et al., 2012).

- *Learnability*: Re-entrance is one of the primary mechanisms by which the hearer can very efficiently learn the meaning of unknown words or constructions. Assume that the hearer has been able to derive the meaning of a sentence, possibly after additional feedback from the speaker, despite the fact that there were items that were not yet part of his conceptual or linguistic inventory. The hearer can then re-enter this meaning in his own production system, taking into account as many constraints derived from the previous sentence as possible, and that way try to come up with a very good guess of what the meaning or function of an unknown item could have been. For example, if the hearer gets the sentence "the block *babodo* the box" and has been able from the context to deduce that a left-of relation holds between the block and the box, then he can infer by abduction that the meaning of the word *babodo* might be 'left-of' (Spranger, 2012).

- *Social Conformity*: Re-entrance is a very effective way by which speaker and hearer can compare their respective language systems and use of language strategies. The hearer can re-enter the meaning gained from the speaker's sentence into his own language production system to find out how he would have formulated the same meaning and can hence detect differences and update his inventory accordingly. Particularly after a failure or misunderstanding has been repaired through further dialog, the speaker can re-enter the sentence in his own comprehension system to discover where the hearer's knowledge might have deviated from his own and then expand or refine his own language system to take this into account.

All these diagnostics and repairs might help speakers and hearers to arrive at a language system that is better suited to their communicative needs but none of them are guaranteed to have this result. There is unavoidably variation in the population and hence a continuous on-going process of selection remains necessary.

5. The Role of Self-Organization

In biological evolution, the self-enforcing causal loop is based on reproductive success: A variant that is more successful tends to generate more offspring, and

hence its frequency increases in future populations. What is the equivalent in the case of cultural language evolution? The selectionist process discussed in the previous sections can already do some of the work, because if a particular variant had communicative success, it means that it had enough expressive power, did not exert excessive cognitive effort and was learnable. Hence there is a chance that it will be used again, which in turn increases the chance that it will be learned and reused by others (Figure 3). The same happens at the level of strategies. The more the language systems built with particular strategies lead to communicative success, the more these strategies are considered by individual language users to be effective (Figure 4), and hence the more the population will prefer them to build their language systems.

But cultural language evolution can make use of another collective process to speed up selection, namely *self-organization* based on *alignment*. Alignment means that speaker and hearer attempt to bring their language systems and choices for language strategies closer to each other. Abundant psycholinguistic evidence (Pickering & Garrod, 2006; Garrod & Doherty, 1994) shows indeed that within the course of a single dialog, dialog partners unconsciously align very rapidly their language systems at all levels: pragmatic, semantic, syntactic, morphological, phonetic, and even gestural. Although linguistic alignment already takes place in the course of a single dialog, it can be expected to leave a long term effect, and this is the basis for accelerating the self-enforcing causal loop: The more a particular variant is used, the more other language users align their own language systems to this variant, and the more coherence is observed within the population.

How is alignment achieved? This depends on the kind of conceptual or linguistic system concerned, and therefore an approach to alignment has to be an integral part of each language strategy:

1. Both the conceptual system and the linguistic system use a variety of categories, many of which are defined in terms of prototypes. Aligning such categories can be achieved by moving their prototypes closer to the feature values of the situation they are talking about. For example, the prototypes of color categories can be moved closer to the actual data when these categories were used in a successful interaction (Steels & Belpaeme, 2005).

2. All lexical and grammatical constructions can be *scored* in the memory of language users. Constructions with the highest scores are chosen by speakers when deciding what and how to say something, and by hearers to interpret what they are hearing. The scores of constructions effectively and success-

fully used within a dialog are increased so that their frequency of use increases and competing variants are inhibited (Steels, 1995).

3. Language users can actively build up dependency networks between constructions based on actual use within the dialogs they engage in so that the successful triggering of one constructions can *prime* other constructions which are thus retrieved faster and are therefore more likely to be used and re-used (Wellens, 2012).

4. *Routinizing* means that a group of components which have been used together in a specific communicative interaction (for example, a network of cognitive functions for conceptualization or a network of grammatical constructions for parsing or producing a complex phrase) are stored as a single unit or chunk, so that they can be retrieved and applied much more efficiently and hence are more likely to be re-used.

Alignment takes place only between the partners in dialog, so the question can be raised whether this is enough to see the population converge towards the same norm. This is where the theory of self-organization provides an important clue. Self-organization is a property of a certain class of distributed systems, namely, systems with a multitude of elements which each autonomously behave in a particular way, for example a set of molecules, ants, or money traders. Although the elements behave autonomously, they are nevertheless influenced by external factors, which are themselves influenced by the behavior of the other elements in the system. Consequently, a positive feedback loop may arise, enforcing a particular variation and locking all elements into the same behavioral pattern. Self-organization is already observed at the level of physics and chemistry, for example in magnets in which spins appear in random directions but align at a critical temperature, or in chemical reactions in which temperature rolls or spiral patterns may appear (Prigogine & Nicolis, 1977).

Self-organization is exploited abundantly in biological systems. The paradigmatic example is path formation in some types of ant societies (Detrain & Deneubourg, 2008). The ants move randomly, but they leave pheromone when returning from a food source to the nest. This pheromone influences the direction of movement of the ants, creating a positive feedback loop between path and ant behavior, and causing ants progressively to form a single path, efficiently transporting food towards the nest until the food source is depleted. Other biological examples of self-organization are bird clouds, termite nests, and patterns on shells or animal skin (Camazine et al., 2001).

Termite nests or beaver dams are the result of a special form of self-organization known as *stigmergy* (Theraulaz and Bonabeau, 1999). In the case of stigmergy, the different elements of the system influence each other through an external structure which they build themselves, namely the nest or dam under construction. In contrast, bird cloud formation is solely based on the local interactions between the birds and is therefore similar to the example of the magnet. Another special form of self-organization is *structural coupling* (Maturana, 2002), which happens when the elements of two *different* systems influence each other's behavior and thus progressively coordinate each other without prior or central coordination.

Self-organization is a general mechanism that occurs in all sorts of natural systems, and it is therefore not surprising that it can help explain how a language community familiar with a set of variants converges on a shared norm, in other words how the language users within a community collectively decide how to conceptualize reality for language and how to express these conceptualizations. Conceptual and linguistic alignment has the same effect as alignment of spins in a magnet or alignment of birds in flight, creating a bird cloud. And it has the same self-enforcing positive feedback loop: The more language users are aligned, the more they converge and become even more aligned.

This crucial role of self-organization is another dimension in which language evolution is different from genetic evolution. In genetic evolution, coherence is mostly reached because genes propagate and spread through vertical transmission. In the social coordination of language based on alignment, convergence is still partly based on transmission because individuals acquire language variants from others, but more importantly on the fact that everybody continuously adapts their language systems and strategies to be similar to those of others. It is true of course that there are significant differences in the attitudes of speakers, with some aligning much more than others. These differences have been shown in psycholinguistic experiments to be an important factor for explaining how far partners are able to establish new communication conventions from scratch (Galantucci, 2005) and it also explains why we may get regional dialects or social strata in language, even if the speakers from these different groups interact frequently with each other as well.

Language has some of the characteristics of stygmergy because written languages or recorded spoken languages are 'left behind' so that other language users can align to them. Utterances can be broadcast to larger audiences and thus influence a larger group of users. We also see structural coupling. The words and grammatical constructions in a language require a particular way of conceptualizing reality, and hence hearers have to align their conceptualizations to those of

speakers in order to be successful in interpreting these utterances. How can they do so if there is no direct experience of these conceptualizations? Structural coupling can explain this phenomenon. Speakers and hearers indirectly get feedback on how far they share conceptualizations based on the outcome of using these conceptualizations in a concrete interaction. Consequently, we see a co-evolution of the conceptual frameworks underlying a language and the way these conceptualizations are expressed.

Computational experiments starting from the mid-nineties (Steels, 1995) have abundantly shown by now that alignment leads to the self-organization of a communication system, similar to the way a path forms in an ant society, and that structural coupling can lead to the coordination of both language and the conceptualizations expressed by language (Steels, 1997). Its importance will be further demonstrated in the remaining chapters of this book, as all experiments discussed later rely on these principles. Moreover the 'semiotic dynamics' observed in human language change, such as the S-shaped curve, can be shown to follow from the self-organization dynamics, and the same mathematical tools used to study self-organization in natural systems like magnets or ant paths can be used to investigate self-organization in language (Castellano et al., 2009).

6. Evolutionary Language Games

What methodology could we use to put more flesh on the theory of cultural language evolution outlined in the previous section?

6.1. Paradigms for studying cultural language evolution

There are three possible paradigms we could use. The first paradigm takes primarily a *linguistic* point of view. It starts from concrete data of language change as found in the historical record or in situations where special circumstances have lead a community of people to develop a new language. This happened for example in the origins of creoles, formed when speakers from diverse language backgrounds suddenly come together and have only weak access to a model language (Mufwene, 2001; Senghas & Coppola, 2001). In such cases, we see lexicalization and grammaticalization processes occurring that can also be observed, but on a slower time scale, in the historical development of 'mature' languages (Traugott & Heine, 1991; Heine et al., 1991). Based on both kinds of data, the kind of cognitive operations and strategies that underly the observed grammaticalization processes can be reconstructed and tested against empirical data (Heine, 1997).

The second paradigm follows a *psychological* approach. It proposes to perform 'semiotic experiments' with human subjects in order to find out what kind of strategies they employ for coming up with a new communication system (Galantucci & Garrod, 2010). These experiments typically put humans in challenging situations where they have to interact without being able to use their existing language. Remarkably, they are able to build up new communication systems rather quickly, even though some people are much better than others (Galantucci, 2005). A special case of such experiments focuses on cultural transmission by creating a 'Chinese whispers' tutor-learner chain, in which the first player acts as tutor for a learner who becomes the tutor for the next learner in the chain (Scott-Phillips & Kirby, 2010).

There is yet another possible paradigm, which is based on *modeling* and this is the one that is explored in the rest of this book. A particular proposal is operationalized and then used to simulate the emergence of language systems in populations of artificial agents. This approach started in the early-nineties (see in particular (Hurford, 1989) and an early review in (Steels, 1998)) and has flourished considerably during the past decade (Briscoe, 2002; Cangelosi & Parisi, 2002; Lyon et al., 2007; Minett & Wang, 2005; Nolfi & Miroli, 2010). The language systems that emerge in these computational experiments are of course never equal to English or Hindi, given the historical contingencies that play a role in normal cultural language evolution, however, by using strategies reconstructed from human languages or by scaffolding the experiment with a vocabulary or partial grammar from an existing human language, the artificial languages are closer to a human source language, which makes the experiment more relevant and the evolution easier to follow.

Even if one chooses the modeling route, there are still many different ways to model cultural language evolution, depending on the specific theory one wants to explore. In the theory of language evolution outlined in the previous section, situated communications between embodied individuals plays a major role in shaping, selecting and self-organizing language systems. We therefore need to employ a modeling approach which has communication (and not only vertical transmission) at its core. Concretely, we will frame communication in terms of *language games*, following up on proposals originally made by Wittgenstein (1953). Language games have also been used in the logic literature to develop a formal game-theoretic semantics of predicates, logical connectives and quantifiers (Hintikaa, 1979), however, the emphasis here is quite different, namely on grounding and embodiment, language used by populations, and above all on-going change and evolution.

The language game approach to the modeling of cultural language evolution started from timid beginnings in the early nineties but right now dozens of exper-

iments have been performed for many different aspects of language, ranging from perceptually grounded vocabularies to grammar. Also there is now a large body of theoretical investigations in the semiotic dynamics that arises in language game experiments. These theoretical investigations are not an add-on luxury, they are crucial to discover what the right system parameters are for language strategies and what the limitations are of chosen methods. The remainder of this section introduces first the notion of a language game as it is used in language evolution research and then surveys the kinds of experiments that are typically done.

6.2. Language Games

A language game is embedded in a cooperative activity in which communication is useful. It attempts to model situated dialog in contrast to the isolated sentences that are commonly used today in formal linguistics. Consequently language game introduce a population of individuals (instead of an idealized speaker) a context, and a communicative purpose, so that pragmatics is part of the modeling effort from the start. Wittgenstein gives the example of the builder and his assistant. The builder requires stones of a certain kind to be handed to him and hence they need a language for talking about types of stones. A language game involves joint attention to some objects and activities in the immediate context and a routinized turn taking interaction that may involve symbolic communication as well as physical actions or gestural communications. A language game takes place based on a specific embodiment that grounds the participants in the world and within a particular environment and ecological setting. These factors co-determine what kind of communicative goals participants may have and what kind of concepts they might be able to use. For example, if the environment contains only black and white objects or if speakers and hearers are all color blind, a hue-based color language cannot (and should not) emerge. If the world is such that objects do not move, a language for talking about events and their temporal structure is irrelevant.

Here is the scenario of a typical language game called the *Naming Game* (see the next chapter for more details (Steels & Loetzsch, 2009)). The Naming Game is a *Game Of Reference*, the speaker attempts to draw the attention of the hearer to an object in the world by naming a characteristic feature of the object. If the object is a specific recognizable individual, then a proper name can be used. It is also possible to name colors, shapes, sizes, as long as they are distinctive.

The game is played by a population `P` of agents and involves a world `W` consisting of objects. Each object is characterized by a point in an n-dimensional feature space,

also called a conceptual space (Gärdenfors, 2000). For example, the color of an object is a point in the three-dimensional color feature space with the dimensions red-green, yellow-blue, and lightness. Two members are randomly selected from the population to take on the roles of speaker and hearer respectively. A context `C` is established which contains a subset of the objects in the world `W`. Then the following interaction takes place:

1. The speaker selects one object out of the context, further called the topic `T`.
2. The speaker finds the distinctive category for the object and names this category.
3. The hearer looks up which object is associated with this category in his memory and examines the context to find out whether there is an object which has this distinctive characteristic.
4. The hearer then signals to the speaker which object was intended according to him, for example by pointing.
5. The speaker checks whether the hearer selected the same object as the one he had originally chosen.
   (a) If they are the same, the game is a success, and the speaker signals this outcome to the hearer.
   (b) If they are different, the game is a failure. The speaker signals this outcome and then points to the topic he had originally chosen.

A 'solution' to the game is a particular language strategy that agents can use to build up a shared set of distinctive categories and names for these categories such that they are successful in the game. The agents do not know these categories nor their names in advance. The language strategy contains diagnostics and repairs for concept formation and concept acquisition and routines for concept alignment, as well as diagnostics and repairs for vocabulary formation and vocabulary acquisition and routines for vocabulary alignment.

There are always many language strategies possible for a language game depending on the specific cognitive functions that are used for playing the game, for learning an existing language system or forming one and particularly for alignment. Each of these strategies has different performance characteristics which can be systematically investigated for the same experimental parameters in order to find the

'linguistic niche' of a strategy. For example, for the Naming Game, we can change the number of objects in the context, the relevant categorial dimensions, how close objects are within the feature space used to form categories, the size of the population, whether the world is dynamic or static, whether the population is dynamic or static, and so on.

Another class of language games are *Action Games*. The speaker tries to get the hearer to do a particular action, such as turn around, raise the left arm, pick up an object or go to a particular location in the room. Action games are particularly useful for studying how names for actions can emerge in a population. One type of Action Game is the *Posture Game* where the speaker does not describe the action but only the bodily posture that he expects the hearer to adopt, such as "arms raised" or "sitting" (Steels & Spranger, 2012).

The Posture Game is again played by a population `P` of agents which have a physical body which they can control to execute actions in the world and a sensory system to get feedback about their own actions (proprioception) and to observe actions by others (through vision). Two members are randomly selected from the population to take on the roles of speaker and hearer respectively.

1. The speaker chooses a posture from his inventory of postures.
2. The speaker retrieves the name for this posture in his vocabulary and transmits that to the hearer.
3. The hearer retrieves the posture by looking up the name in his own vocabulary and evokes the motor behavior that could achieve this posture.
4. The speaker observes the posture adopted by the hearer and checks whether it fits with the prototypical visual body-image of the posture he had originally chosen.
   (a) If this is not the case, the speaker signals failure. The speaker activates his own motor behavior for achieving this posture in order to repair the communication, so that there is an opportunity for the hearer to learn the speaker's name for this posture.
   (b) Otherwise the speaker signals success.

Again, this game definition is just a setting. The solution takes the form of concrete proposals for language strategies by which speaker and hearer can invent, learn, and coordinate names for postures, as well as learn the visual image schemata of

a posture, the motor control programs to achieve the posture, and the associations between the two. Language games almost always raise many fundamental issues in cognitive science. For example, to be able to play the posture game, the players need to have a mirror system so that they can recognize actions of others in terms of their own actions (Rizzolatti & Arbib, 1998). But there are also recurrent problems that come up in almost every game. For example, both the Naming Game and the Action Game require that the population establishes lexical conventions, even though the game script, the conceptual system, and the strategies of the agents are different.

It is almost impossible to understand the consequences of a proposed language strategy, particularly for complex strategies with rich conceptualizations and complex grammar, which is the main reason why computer simulations and robotic experiments of language games are necessary. A language game experiment is intended to test an evolutionary explanation for some aspect of language, for example, why languages have 'fuzzy' quantifiers, such as "some", "a few", "almost all", etc. (Pauw & Hilferty, 2012), so the first step is always to be very clear about the phenomena one wants to investigate.

Setting up an experiment requires defining (i) an environment and an embodiment of the agents, (ii) a communicative goal within this environment, which is assumed to be part of a larger cooperative interaction that is usually not modeled, (iii) a game script, that not only deals with routine use but also with feedback and repair, and (iv) a population structure. The challenge is then to come up with an operational language strategy so that agents endowed with this strategy cannot only routinely use the aspect of language under study but also learn it from others and invent and coordinate a language system from scratch given this strategy. Some experiments go one step further and provide agents only with the necessary components to evolve new language strategies and try them out in a selectionist semiotic dynamics, but the goal of the experiment is the same, namely to arrive at a shared communication system that is adequate for the embodiment, environment, and cooperative goal of the language game.

If carried out properly, language game experiments are able to establish in an objective fashion whether a proposed strategy works, what its performance characteristics are, how different alternative strategies compare and might compete with each other in a selectionist dynamics and how new strategies could form by recruitment or by exaptation from an existing strategy.

6.3. Empirical Relevance of Language Game Experiments

Artificial language evolution experiments have a number of advantages compared to psychological experiments or linguistic reconstructions: We can monitor the complete state of all the artificial agents as the evolutionary process unfolds. We can take out or add components to language strategies and see their effect. We can roll back the clock and start an experiment again to see how alternative solutions are explored. And we can systematically compare different theories of language evolution by testing how well they are able to cope with the communicative challenges arising in a particular language game. Of course, to tackle significant properties of human language, the agents need to become more complex because language itself is complex and it is about a complex world, but the state of the art in computer science, Artificial Intelligence, computational linguistics, and robotics is now far enough advanced that artificial language evolution experiments are entirely feasible, as the many examples contained further in this book show.

The sceptical cognitive scientist or neuroscientist tends to object that the operationalization of the cognitive functions needed by language strategies are not based on models of human neurobiological structures and processes and that the mechanisms used by the human brain might therefore be very different. This is of course true, but, first of all, at least some of the processes responsable for language evolution are of a cultural nature and do not depend on the particular characteristics of human brain implementation. Just like one can study economic or social decision-making without necessarily having a detailed model of the neurobiological processes involved, we can study some of the principles relevant for the origins and development of language and cognition by staying at a functional and cultural level.

Second, the functional models we are using have an important advantage over the boxology of classical neurology because each component in the model has been operationalized and it has been shown that the whole system is indeed able to perform the functions that are ascribed to it. Hence the models could help to elucidate better which kind of information processing the brain needs to carry out in principle. The wing of an airplane is not made out of feathers but it has the same function as a bird wing and the construction and investigation of artificial wings therefore tells us a lot about natural wings.

The sceptical linguist tends to object that the artificial languages emerging in these experiments are not exactly like human languages, indeed, they can never be, and that therefore these experiments are merely theoretical. This criticism is valuable but can be addressed by starting the investigation how a particular class of

linguistic phenomena (for example tense-aspect systems) might arise by a *reconstruction* of a language system that is actually observed in a human language. This suggests that an ideal research project in evolutionary linguistics should go through the following steps:

1. The first step is to *select* a particular language strategy and an example language where this strategy is clearly used. For example:

- We might be interested in spatial language, and specifically in the expression of spatial categories with lexical items, as in "left" vs. "right" or "near" vs. "far" (Spranger, 2012).

- We might be interested in a graded color term strategy, as in "*very* blue" or "*slightly* blue", and use Tarahumara, a North Mexican indiginous language, as example language because it has a system of morphological affixes to express graded color terms (Bleys, 2012).

- We might be interested in the expression of epistemic modality, which expresses the speaker's attitude with respect to the truth, probability, certainty, belief or evidence of propositions, and use Spanish as an example language because it has a prominent language system for mood and modality (Beuls, 2012).

Often an adequate description of the language system can already be found in the linguistic literature, and a representative corpus can be found or constructed with example sentences for later testing.

2. The second step is to understand the *ecological or functional significance* of the language system: What communicative need is being addressed? Why is this system there? This analysis has a direct impact on the design of a language game where this communicative need arises. The game and the environment of the agents should be such that the communicative need is indeed present and therefore acts as a trigger to stimulate concept formation as well as lexicalization and grammaticalization processes. For example:

- To study spatial language, a game of reference could be used where the objects are not identifiable through proper names or descriptions of their visual characteristics (such as color) only, so that it becomes relevant and even necessary to use the position of objects with respect to the partners in the game or with respect to other objects acting as landmarks. This approach is used in the later chapter by Spranger (2012).

- To study case strategies, a description game could be used in which the speaker describes the event occurring in the scene and the hearer agrees or disagrees that the description is valid for the scene. Often it is necessary to describe the roles of the participants in an event because the same type of event may occur more than once, and this triggers the need for expressing argument structure (van Trijp, 2012).

- Part of grammatical structure is about avoiding the cognitive effort of the hearer in parsing (Steels & Wellens, 2006), and so a language game challenge can be set up that involves hierarchical constituent structure which is not explicitly marked. This then acts as a framework for studying the emergence of constituent ordering or of an agreement system, as shown in the later chapter by Beuls et al. (2012).

3. The third step is to *reconstruct the language system* that was chosen as case study. This means: to implement representations for the meanings to be expressed and for the processes that ground these meanings in reality through the sensorymotor system of the agent, and to reconstruct the lexical and grammatical constructions for parsing and producing sentences. A *reconstruction experiment* should establish with a corpus of utterances or with an experimental data set obtained from artificial robotic agents a performance baseline. Usually a particular language system does not lead to 100 % communicative success, because in normal sentences and natural dialog, several language strategies cooperate. For example, basic color terms may not be accurate enough to identify a color sample, in which case a combination of basic colors with graded membership terms (such as "slightly blue") becomes useful. Aspects of language which do not involve the target language system under investigation but are nevertheless necessary to get produce or comprehend sentences can be scaffolded. For example, in grammatical experiments it is often useful to scaffold an initial vocabulary.

4. The fourth step is to *operationalize the language strategy*. This means: to implement operators for inventing, acquiring, and aligning both the conceptual system and the linguistic system covered by this strategy. The adequacy of the acquisition and alignment operators can be objectively tested through an *acquisition experiment* whereby some agents in the population act as accomplished speakers (based on the results of the reconstruction in step 3) and others act as learners that use an implementation of the operational language strategy but have no prior knowledge of the language system.

5. The next step is to perform a *language formation experiment* in which none of the agents is given language systems, but all of them are endowed with the same language strategy. The agents then engage in a series of language games where they invent, acquire, and align their private language systems in the hope that a shared language system emerges. The emergent language system will of course not be identical to the one found in the target language, but it should have strong similarities. For example, a basic set of color categories and color terms emerges, but it may not be the one found in English or Polish, or an agreement system emerges but it the semantic or syntactic features that have to agree (for example person and number) are not the ones found in German or Hungarian. The success in reaching this step can again be objectively evaluated, by tracking communicative success, convergence and the relation to the naturally occurring language system that was used as the point of departure.

6. The final step is to try and explain how the language strategy itself could emerge. In this case, the experimenters take one further step back and endow agents only with a pool of possible cognitive functions that they can recruit and configure to build the conceptualization and language processing needed to enact a language system and the invention, adoption and alignment operators that allow them to self-organization a system from scratch. The success of a self-assembled strategy in a *strategy formation experiment* depends on the success the agents have with the language systems they were able to build and coordinate using this strategy. There is now unavoidable variation, both in terms of the paradigmatic choices for the language system and for the possible language strategies, so that semiotic dynamics takes place at two levels: the level of language strategies, which are now competing or cooperating with each other, and the level of language systems, where different paradigmatic choices need to arise and become coordinated. This kind of semiotic dynamics is explored in *strategy competition experiments*.

7. Conclusions

If we accept that language is not only a biological and social phenomenon but also a cultural phenomenon that arises by the cognitive activities of human language users, then an explanation for the origins and evolution of language requires that we identify which cognitive functions and interaction patterns are needed for the emergence of the kind of conceptual systems and linguistic systems we find in human natural languages and how these functions can give rise to shared language systems based on shared language strategies. This chapter introduces briefly a theory for cultural language evolution based on the notions of selection and self-organization.

Selection means that there is a distinct process of generating variants and of testing variants and a self-enforcing causal loop that ensures that variants which satisfy the desired selection criteria thrive and become dominant in the population. Linguistic self-organization is based on alignment. Speakers and hearers align their language systems and strategies at all levels and this will lead to convergence without prior design or central coordination.

Of course, this theory outlined in this chapter now needs to be instantiated by investigating concrete language systems and language strategies and exercising them through evolutionary language game experiments. This task is certainly not easy because it presupposes worked out theories of the cognitive functions needed for language and theories about language processing and language learning. In each case study, we will need to answer the following questions:

1. What is the functional niche of this language system, what communicative purpose does it serve?

2. What are the cognitive functions needed to deal with the grounded semantics of the conceptual system involved?

3. What are the linguistic design patterns that express the conceptualizations produced by a conceptual system, e.g. through individual words such as basic color terms, through simple item-based constructions, or through more elaborate syntactic systems, like internal agreement relations within a constituent?

4. What are the diagnostics and repairs that progressively build up a language system so that the different linguistic selection criteria (expressive adequacy, minimal cognitive effort, learnability and social conformity) have a higher chance to be satisfied?

5. What kind of alignment could drive the population towards conceptual and linguistic convergence so that the self-enforcing causal loop between the outcome of communication and the paradigmatic choices within a language system or the widespread use of a strategy can be established?

Subsequent chapters of this book show that answering these questions is non-trivial but not impossible. Different language strategies or language systems are going to be studied, for example, how a system of determiners could arise (Pauw & Hilferty, 2012), how the emergence of action language can be explained (Steels &

Spranger, 2012), how argument structure could become expressed in a case system (van Trijp, 2012), how an agreement system within nominal phrases might arise (Beuls et al., 2012) and so on. Each of these case studies provides increasing evidence that the fundamental principles of linguistic selection and self-organization provide a very productive paradigm for explaining the rich complex structures we find in human languages.

Acknowledgements

Research for this chapter was carried out at the Artificial Intelligence Laboratory of the University of Brussels (VUB) and the Sony Computer Science Laboratory in Paris. The author thanks all members of the language groups at these laboratories for their highly valued input in working out the ideas expressed in this paper. Remi van Trijp, Bart de Boer and anonymous reviewers provided valuable comments to improve the present paper.

References

Aitchinson, Jane (2001). *Language Change Progress or Decay?* Cambridge UK: Cambridge University Press.

Arthur, Brian (Ed.) (1996). *The Economy as a Complex Adaptive System*. Santa Fe Institute Series on Complexity. Menlo Park, Ca: Addison-Wesley.

Bailey, Charles (1973). *Variation and Linguistic Theory.* Washington: Center for Applied Linguistics.

Beuls, Katrien (2012). Spanish modals in Fluid Construction Grammar. In Luc Steels (Ed.), *Computational Issues in Fluid Construction Grammar*. Berlin: Springer Verlag.

Beuls, Katrien, Luc Steels, Sebastian Höfer (2012). The emergence of internal agreement systems. In Luc Steels (Ed.), *Experiments in Cultural Language Evolution*. Amsterdam: John Benjamins.

Bickerton, Derek, Eors Szathmary (2008). *Biological Foundations and Origin of Syntax. Strungmann Forum Reports, vol. 3.* Cambridge, MA: The MIT Press.

Bleys, Joris (2012). Language strategies for color. In Luc Steels (Ed.), *Experiments in Cultural Language Evolution*. Amsterdam: John Benjamins.

Boyd, Richard, Peter Richerson (2005). *The Origins and Evolution of cultures*. Oxford: Oxford University Press.

Briscoe, Ted (Ed.) (2002). *Linguistic Evolution through Language Acquisition: Formal and Computational Models*. Cambridge: Cambridge University Press.

Camazine, Scott, Jean-Louis Deneubourg, Nigel Franks, James Sneyd, Guy Theraulaz, Eric Bonabeau (2001). *Self-Organization in Biological Systems*. Princeton: Princeton University Press.

Cangelosi, Angelo, Domenico Parisi (2002). Computer simulation: A new scientific approach to the study of language evolution. In Angelo Cangelosi, Domenico Parisi (Eds.), *Simulating the Evolution of Language*, 3–28. London: Springer Verlag.

Castellano, Claudio, Santo Fortunato, Vittorio Loreto (2009). Statistical physics of social dynamics. *Review of Modern Physics*, 81, 591 – 646.

Cavalli-Sforza, Luigi, Marcus Feldman (1981). *Cultural transmission and evolution: a quantitative approach*. Princeton, NJ, USA: Princeton University Press.

Christensen, Clayton (1997). *The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail.* Cambridge Ma: Harvard Business School Press.

Croft, William (2000). *Explaining Language Change: An Evolutionary Approach*. Harlow Essex: Longman.

Croft, William (2001). *Radical Construction Grammar*. Oxford: Oxford University Press.

Dawkins, Richard (1976). *The Selfish Gene*. Oxford: Oxford University Press.

Detrain, Claire, Jean-Louis Deneubourg (2008). Collective decision-making and foraging patterns in ants and honeybees. *Advances in Insect Physiology*, 35, 123–173.

Diessel, Holger (1999). *Demonstratives: Form, Function, and Grammaticalization.* Amsterdam: John Benjamins Pub. Co.

Eco, Umberto (2000). *The Search for the Perfect Language.* New Haven: Lawrence Erlbaum Ass.

Edelman, Gerald (1987). *Neural Darwinism. The theory of neuronal group selection.* New York: Basic Books.

Galantucci, Bruno (2005). An experimental study of the emergence of human communication systems. *Cognitive Science*, 29(5), 737–767.

Galantucci, Bruno, Simon Garrod (2010). Experimental semiotics: A new approach for studying the emergence and the evolution of human communication. *Interaction Studies*, 11(1), 1–13.

Gärdenfors, Peter (2000). *Conceptual Spaces: The Geometry of Thought*. A Bradford Book. MIT Press.

Garrod, Simon, Gwyneth Doherty (1994). Conversation, co-ordination and convention: an empirical investigation of how groups establish linguistic conventions. *Cognition*, 53(3), 181–215.

Garrod, Simon, Nicolas Fay, Shane Rogers, Bradley Walker, Nik Swoboda (2010). Can iterated learning explain the emergence of graphical symbols? *Interaction Studies*, 11(1), 52–69.

Goldberg, Adele (2003). Constructions: a new theoretical approach to language. *Trends in Cognitive Sciences*, 7(5), 219–224.

Heine, Bernd (1997). *The Cognitive Foundations of Grammar*. Oxford: Oxford University Press.

Heine, Bernd, Ulrike Claudi, Friederike Hünnemeyer (1991). *Grammaticalization: A Conceptual Framework*. Chicago: University of Chicago Press.

Heine, Bernd, Tania Kuteva (2002). *World Lexicon of Grammaticalization*. Cambridge University Press, Cambridge, MA.

Heine, Bernd, Tania Kuteva (2008). *The Genesis of Grammar. A reconstruction.* Oxford: Oxford University Press.

Hintikaa, Jaako (1979). *Game-theoretical Semantics*. Dordrecht: Reidel Pub. Co.

Holland, John (1992). Complex adaptive systems. *Daedalus*, 121(1), 17–30.

Hurford, James (1989). Biological evolution of the saussurean sign as a componnet of the language acquisition device. *Lingua*, 77, 187–222.

Jablonka, Eva, Marion Lamb (2005). *Evolution in Four Dimensions. Genetic, Epigenetic, Behavioral, and Symbolic Variation in the History of Life*. Cambridge Ma: The MIT Press.

Kay, Paul, Charles Fillmore (1996). Grammatical constructions and linguistic generalizations: the what's x doing y? *Language*, 72, 1–33.

Knight, Chris, James Hurford, Michael Studdert-Kennedy (Eds.) (2000). *The Evolutionary Emergence of Language: Social Function and the Origins of Linguistic Form*. Cambridge: Cambridge University Press.

Knight, Chris, Camilla Power, Ian Watts (2001). The human symbolic revolution: A darwinian account. *Cambridge Archaeological Journal*, 5(1), 75–114.

Koza, John (1992). *Genetic programming: On the programming of computers by means of natural selection*. Cambridge, MA: The MIT Press.

Levelt, Willem (1989). *Speaking*. Cambridge MA: MIT Press.

Luisi, Luigi, Thomas Oberholzer (2001). Origin of life on earth: Molecular biology as an approach to the minimal cell. In *The Bridge between the Big Bang and Biology.*, 345–355. Rome: CNR Press.

Lyon, Caroline, Chrystopher Nehaniv, Angelo Cangelosi (Eds.) (2007). *Emergence of Language and Communication*. Lecture Notes in Computer Science. Berlin: Springer Verlag.

MacLaury, Robert (1992). From brightness to hue: an explanatory model of color-category evolution. *Current Antropology*, 33(2), 137–187.

Maynard Smith, John, Eors Száthmary (1993). The origin of chromosomes i. selection for linkage. *Journal of Theoretical Biology*, 163, 437–446.

Michaelis, Laura, Knud Lambrecht (1996). Toward a construction-based theory of language function: The case of nominal extraposition. *Language*, 72, 215–247.

Minett, John, William Wang (2005). *Language Acquisition, Change and Emergence: Essays in Evolutionary Linguistics*. Hong Kong: City University of Hong Kong Press.

Mufwene, Salikoko (2001). Competition and selection in language evolution. *Selection*, 3(1).

Nolfi, Stefano, Dario Floreano (2000). *Evolutionary Robotics*. Cambridge Ma.: MIT Press.

Nolfi, Stefano, Marco Miroli (Eds.) (2010). *Evolution of Communication and Language in Embodied Agents*. Berlin: Springer Verlag.

Pauw, Simon, Joe Hilferty (2012). The emergence of quantification. In Luc Steels (Ed.), *Experiments in Cultural Language Evolution*. Amsterdam: John Benjamins.

Pickering, Michael, Simon Garrod (2006). Alignment as the basis for successful communication. *Research on Language and Computation*, 4, 203–228.

Prigogine, Ilya, Gregoire Nicolis (1977). *Self-Organization In Non-Equilibrium Systems.* New York: wiley.

Rizzolatti, Giacomo, Michael Arbib (1998). Language within our grasp. *Trends in Neurosciences*, 21(5).

Scott-Phillips, Thomas, Simon Kirby (2010). Language evolution in the laboratory. *Trends in Cognitive Sciences*, 9(14), 411–417.

Senghas, Ann, Marie Coppola (2001). Children creating language: Jow nicaraguan sign language acquired a spatial grammar. *Psychological Science*, 12(4), 537–558.

Sperber, Dan, Deirdre Wilson (1986). *Relevance: Communication and Cognition*. Cambridge, MA: Harvard University Press.

Spranger, Michael (2012). The co-evolution of basic spatial terms and categories. In Luc Steels (Ed.), *Experiments in Cultural Language Evolution*. Amsterdam: John Benjamins Co.

Steels, Luc (1995). A self-organizing spatial vocabulary. *Artificial Life*, 2(3), 319–332.

Steels, Luc (1997). Constructing and sharing perceptual distinctions. In Maarten W. van Someren, Gerhard Widmer (Eds.), *Proceedings of the European Conference on Machine Learning*, 4–13. Prague, Czech Republic: Springer.

Steels, Luc (1998). Synthesising the origins of language and meaning using co-evolution, self-organisation and level formation. In James Hurford, Michael Studdert-Kennedy, Chris Knight (Eds.), *Approaches to the Evolution of Language: Social and Cognitive Bases*, 384–404. Edinburgh: Edinburgh University Press.

Steels, Luc (2000). Language as a complex adaptive system. In Marc Schoenauer (Ed.), *Proceedings of PPSN VI*, Lecture Notes in Computer Science, 17–26. Berlin: Springer-Verlag.

Steels, Luc (2003). Language re-entrance and the 'inner voice'. *Journal of Consciousness Studies*, 10(4-5), 173–185.

Steels, Luc (2007). The recruitment theory of language origins. In Caroline Lyon, Chistopher Nehaniv, Angelosi Cangelosi (Eds.), *The Emergence of Communication and Language*, 129–151. Berlin: Springer Verlag.

Steels, Luc (Ed.) (2012). *Design Patterns in Fluid Construction Grammar*. Amsterdam: John Benjamins Pub.

Steels, Luc, Tony Belpaeme (2005). Coordinating perceptually grounded categories through language: A case study for colour. *Behavioral and Brain Sciences*, 28(4), 469–489.

Steels, Luc, Martin Loetzsch (2009). Perspective alignment in spatial language. In Kenny R. Coventry, Thora Tenbrink, John. A Bateman (Eds.), *Spatial Language and Dialogue*, 70–89. Oxford University Press.

Steels, Luc, Michael Spranger (2012). Emergent mirror systems for body language. In Luc Steels (Ed.), *Experiments in Cultural Language Evolution*. Amsterdam: John Benjamins.

Steels, Luc, Remi van Trijp (2012). How to make construction grammars fluid and robust. In Luc Steels (Ed.), *Design Patterns in Fluid Construction Grammar*. Amsterdam: John Benjamins.

Steels, Luc, Pieter Wellens (2006). How grammar emerges to dampen combinatorial search in parsing. In Paul Vogt, Yuuya Sugita, Elio Tuci, Chrystopher Nehaniv (Eds.), *Symbol Grounding and Beyond. Proceedings of the Third EELC*, LNAI 4211, 76–88. Berlin: Springer-Verlag.

Thomason, Sarah, Terence Kaufman (1988). *Language Contact, Creolization, and Genetic Linguistics*. Berkeley: University of California Press.

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

Traugott, Elisabeth, Bernd Heine (Eds.) (1991). *Approaches to Grammaticalization*, *Typological Studies in Language 19*, vol. 1. Amsterdam: John Benjamins.

Van Kemenade, Ans (1987). *Syntactic Case and Morphological Case in the History of English.* Dordrecht: Forist Publications.

van Trijp, Remi (2012). The evolution of case systems for marking event structure. In Luc Steels (Ed.), *Experiments in Cultural Language Evolution*. Amsterdam: John Benjamins.

Wellens, Pieter (2012). Organizing constructions in networks. In Luc Steels (Ed.), *Design Patterns in Fluid Construction Grammar*. Amsterdam: John Benjamins.

Wellens, Pieter, Martin Loetzsch, Luc Steels (2008). Flexible word meaning in embodied agents. *Connection Science*, 20(2), 173–191.

Wittgenstein, Ludwig (1953). *Philosophical Investigations*. New York: Macmillan.