1   Introduction

1.1   Preliminary Remarks

There has been a long tradition in linguistics and related disciplines to separate the study of the linguistic system (i.e., langue, competence) from the study of language use (i.e., parole, performance). In formal linguistic theory, grammar is a self-contained, deductive system consisting of discrete categories and algorithmic rules that are usually analyzed without any consideration of how language is used and processed (e.g., Chomsky 1965, 1986).

This view of grammar has been challenged, however, by usage-based linguists and psychologists who have argued that linguistic knowledge, including knowledge of grammar, emerges from language use (e.g., Tomasello 2003; Goldberg 2006; Bybee 2006). In the usage-based approach, grammar is seen as a dynamic system consisting of fluid structures and flexible constraints that are shaped by general mechanisms of communication, memory and processing. Specifically, these researchers claim that grammar constitutes a network that is constantly restructured and reorganized under the influence of domain-general processes of language use (see Diessel 2017 for a review).

In order to understand the dynamics of the grammar network, usage-based researchers study the development of linguistic structure, both in history and language acquisition. One factor that has a great impact on language development is frequency of occurrence. As frequency strengthens the representation of linguistic elements in memory, it facilitates the activation and processing of words, categories and constructions, which in turn can have long-lasting effects on the development of linguistic structure. There is a large body of research indicating that frequency is an important determinant of language use, language acquisition and language change and that the cognitive organization of grammar is crucially influenced by language users' experience with particular lexemes and constructions (e.g., Bybee and Hopper 2001; Ellis 2002; Diessel and Hilpert 2016).

This book provides a comprehensive overview and discussion of usage-based research on grammar and grammatical development. The usage-based approach draws on research in functional and cognitive linguistics (e.g., Croft

1

2001; Hay 2003; Stefanowitsch and Gries 2003; Bybee 2010; Traugott and Trousdale 2013; Perek and Goldberg 2015) and related research in cognitive psychology and cognitive science (e.g., Elman et al. 1996; Seidenberg and MacDonald 1999; Christiansen and Chater 2008; Fedzechkina et al. 2013; Steels 2015; Ellis et al. 2016). These fields of research complement each other, but as it stands they are only loosely connected. It is the purpose of this book to integrate the various strands of research into a more unified framework and to elaborate some central principles of the usage-based approach. In particular, the book sets out to elaborate the network view of grammar.

It is a basic assumption of the usage-based approach that linguistic knowledge is organized in an associative network (e.g., Beckner et al. 2009), but although the network view of language is frequently invoked in the usage-based literature, it has not yet been developed into an explicit theory or model. To be sure, there are network accounts of morphology and lexical semantics, but syntactic phenomena are only rarely analyzed in the framework of a network model (see Diessel 2015 for discussion).

In this book, we will consider a structured network model of grammar in which all aspects of linguistic structure, including core concepts of syntax (e.g., noun, case, subject), are analyzed in terms of associative connections between lexemes, categories and constructions. The model is inspired by computational research with neural networks (e.g., Rumelhart and McClelland 1986a; Elman et al. 1996) and is intended to provide a unified framework for the analysis of language use and linguistic structure (e.g., Bates and MacWhinney 1989; Bybee 2006). Before we consider the details of the model, let us briefly consider three basic principles of the usage-based approach as a background for the subsequent discussion (see Diessel 2011a).

1.2 Three General Principles of Usage-Based Linguistics

The usage-based approach challenges basic principles of linguistic research that have long been taken for granted. In particular, it challenges the conception of three general divisions that have provided the foundation of syntactic theory since the advance of generative linguistics in the 1950s and 1960s, namely, (i) the division between linguistic knowledge and language use, or competence and performance, (ii) the dichotomy of synchronic states and diachronic development and (iii) the distinction between words and rules.

1.2.1 Linguistic Knowledge and Language Use

All (contemporary) linguists conceive of language, notably grammar, as a cognitive system that involves linguistic knowledge, but generative and

usage-based researchers make very different assumptions about the nature and origin of linguistic knowledge and its relationship to language use (Newmeyer 2003; Bybee 2006).

In the classic version of the generative approach, knowledge of grammar is grounded in a particular faculty of the mind including categories and rules, or constraints, that are part of our genetic endowment and that can generate an infinite number of sentences (Chomsky 1986; Pinker and Jackendoff 2005). Language use, or performance, involves grammatical knowledge, commonly referred to as competence, but is also influenced by general psychological processes such as sensory perception and attention that do not immediately concern the representation of grammatical knowledge. Building on this view, generative linguists separate the study of grammar, or competence, from the study of language use, as the latter involves "performance phenomena," caused by general psychological processes, that are not part of the language faculty (and therefore commonly excluded from syntactic theory).

Usage-based researchers reject the innateness hypothesis of generative linguistics and with it the related distinction between competence and performance. In the usage-based approach, language is seen as a "complex adaptive system" that has evolved for the purpose of communication and processing (e.g., Steels 2000; Beckner et al. 2009). Rather than claiming that grammatical concepts are grounded in a particular faculty of the mind, usage-based linguists argue that all aspects of linguistic knowledge, including the core concepts of grammar, emerge from general cognitive mechanisms that are not only involved in the use of language but also in other cognitive phenomena such as vision, memory and decision-making. In accordance with this view, these researchers seek to explain how linguistic structure is shaped by (nonlinguistic) factors of performance, or as Bybee (2010: 1) puts it, it is the general goal of usage-based linguistics "to derive linguistic structure from the application of domain-general processes."

1.2.2   Synchronic States and Language Development

In order to study the (long-term) effects of language use on linguistic knowledge, one has to consider the way in which linguistic structures evolve over time. Ever since Saussure ([1916] 1994), the field of linguistics has been divided into two major research areas: synchronic linguistics, which is concerned with the analysis of linguistic states at a particular point in time, and diachronic linguistics, which is concerned with the analysis of language change. Prior to Saussure, linguistic structure was generally analyzed in light of its development – synchrony and diachrony were studied together in a unified framework (Paul [1880] 1920). But since the advance of linguistic structuralism, the study of synchronic states and language change has been split

into separate fields of research with distinct goals and different methods. The division of labor has been reinforced by the innateness hypotheses of generative grammar. If grammar is grounded in a particular faculty of the mind, language change concerns only the periphery of grammar and the innate core can be studied from a purely synchronic perspective.

Usage-based linguists have questioned the usefulness of the structuralist division between synchronic and diachronic linguistics. If we conceive of grammar as an emergent system, all aspects of linguistic structure, including the core concepts of syntax, are subject to change, and in order to understand the nature of this system, one has to study language development, both in history and acquisition. This explains why usage-based linguists have emphasized the importance of grammaticalization for syntactic theory (Boye and Harder 2012) and why some usage-based scholars have turned to the study of language acquisition (Goldberg 2006). In the structuralist paradigm, grammatical research is primarily concerned with the analysis of linguistic states, but in the usage-based model, the focus of analysis is on the dynamics of the linguistic system (Hopper 1987).

1.2.3   Words and Rules

Finally, usage-based linguists have challenged the traditional distinction between words and rules, which is perhaps the most fundamental dichotomy of (traditional) linguistic theory (Pinker 1999). Words are signs or symbols that combine a particular phonetic form with a particular concept or meaning, whereas rules are commonly defined as (cognitive) algorithms that serve to combine abstract categories into larger structures. Phrase structure rules, for instance, combine word class categories (and phrases) into syntactic constituents ($PP \rightarrow P\ NP$).

On this view, linguistic rules are completely different entities from words or lexemes, which are stored and processed in different modules of the mind. In the classic version of generative grammar, language consists of two general components: the mental lexicon, which includes words and idiomatic expressions, and grammar, which includes syntactic categories and rules or constraints (Chomsky 1965, 1986).

The distinction between lexicon and grammar has been a cornerstone of linguistic theory, but this distinction has lost some of its importance over the past 25 years as an increasing number of theoreticians has argued that linguistic structure is licensed by constructions rather than by algorithmic rules (Fillmore et al. 1988; Goldberg 1995). A construction is a holistic pattern in which a particular configuration of structural elements is associated with a particular function or meaning. A noun phrase such as *John's car*, for instance, can be seen
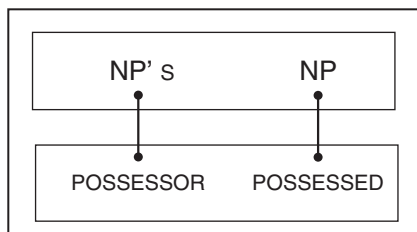
Figure 1.1 The English genitive construction

as a construction (with two slots for nominal expressions) that typically desig-
nates a particular semantic relationship of possessor and possessed (Figure 1.1).

The notion of construction is of central significance to the usage-based
analysis of grammar. In fact, usage-based linguists have drawn so frequently
on theoretical concepts of construction grammar that the two approaches are
often presented as a unified framework (Tomasello 2003; Goldberg 2006).
Note, however, that the notion of construction grammar subsumes a whole
family of related theories that are not all usage-based (see Hoffmann and
Trousdale 2013 for an overview). Indeed, one of the earliest and most influen-
tial construction-based theories, that is, the sign-based theory of construction
grammar developed by Fillmore and Kay (1999), explicitly maintained the
generative conception of competence and performance and paid little attention
to usage and development. However, other varieties of construction grammar
take a dynamic perspective and have made important contributions to the
usage-based approach (e.g., Croft 2001; Goldberg 2006; Steels 2013; Hilpert
2014).

1.3   Goal and Scope of the Book

To summarize the previous discussion, usage-based linguists conceive of
language as a dynamic system of emergent structures and flexible constraints
that are in principle always changing under the pressure of domain-general
processes, that is, processes that do not only concern the use of language but
also nonlinguistic cognitive phenomena such as visual perception, memory
retrieval and automatization. The focus of analysis is on the development of
linguistic structure rather than on particular linguistic states.

The emergentist view of linguistic structure has far-reaching consequences
for the study of grammar. Traditionally, grammatical analysis presupposes
a "toolkit" of primitive categories that are defined prior to the analysis of any
particular structure (Jackendoff 2002: 75). The "toolkit approach" has domi-
nated syntactic theory for many decades (see Croft 2001 for discussion); but if

we think of language as a dynamic system of emergent structures and fluid constraints, we cannot approach the study of grammar with a predefined set of primitive categories. On the contrary, what we need to explain is how linguistic categories evolve, stabilize and change. The underlying hypothesis is that all aspects of linguistic structure, including the most basic categories, such as noun, word and phrase, are emergent and fluid.

That does not mean, however, that linguistic structure is completely unconstrained in the usage-based model of grammar. Like any other grammatical theory, the usage-based theory of grammar rests on particular assumptions about the nature of linguistic structure and the overall organization of the linguistic system. As I see it, there are two general aspects of cognition that constrain grammar in the usage-based approach: (i) the domain-general processes that shape linguistic structure in the process of language use, and (ii) the network architecture of the grammatical system. It is the general goal of this book to elaborate on these two aspects of the usage-based model and to combine them into a coherent account.

The two following chapters introduce the basic assumptions of the current approach. We begin with the architecture of the grammar network (Chapter 2) and then turn to domain-general processes of language use (Chapter 3). In the remainder of the book, we will consider the various aspects of the model in more detail. Each chapter is concerned with particular cognitive processes and a specific aspect of the network model.

Part I

Foundations

2 Grammar as a Network

2.1 Introduction

It is a standard assumption of usage-based linguistics that grammatical structure consists of signs, that is, constructions, that are associated with each other in various ways so that the entire inventory of linguistic signs is commonly characterized as some kind of network. The network view of grammar has been expressed in a large number of studies (e.g., Bybee 1985; Langacker 1988; Bates and MacWhinney 1989; MacDonald et al. 1994; Goldberg 1995; Elman et al. 1996; Pickering and Branigan 1998; Croft 2001; Fried and Östman 2005; Hoey 2005; Hudson 2007; Steels 2011; Brown and Hippisley 2012; Traugott and Trousdale 2013; Hilpert 2014; Schmid 2016; van Trijp 2016; Ellis et al. 2016), but it has not yet been developed into an explicit theory or model (see Diessel 2015 for discussion).

Some researchers have argued that constructions are interconnected by particular types of links. Inheritance links, for instance, define connections between constructions at different levels of abstraction, and polysemy links specify relations between different senses of constructions (Goldberg 1995: 67–100; Hilpert 2014: 50–73). However, while the proposed links are certainly important, they are not sufficient to account for all aspects of linguistic structure. As we will see, in order to explain the full range of phenomena that constitute a speaker's linguistic knowledge, we need to consider additional types of links or relations.

This book delineates a theoretical network model of grammar in which all concepts of grammar (e.g., constituent structure, argument structure, word classes, grammatical relations, morphological paradigms and constructions) are defined by various types of links, or relations, that indicate associations between different linguistic elements. The present chapter introduces the general architecture of the network model, which will then be elaborated in the remainder of the book.

2.2 Some General Properties of Networks

Let us begin with some general remarks on the use of networks in science. Network models are used in many scientific disciplines to analyze a wide range

9

of phenomena, e.g., ecosystems, social relations, the brain, economic circuits, traffic systems, cognitive processes and language (Buchanan 2002). On the face of it, these phenomena seem to be completely unrelated, but the network models that are used to analyze them have some interesting properties in common, which has led some scholars to argue that science needs a general "theory of networks" (Buchanan 2002; see also Baronchelli et al. 2013).

Formally, a network consists of two basic entities: (i) a set of nodes, sometimes referred to as vertices, and (ii) connections, also called arcs, links, relations or edges (Buchanan 2002). One reason why network models are so frequently used across disciplines is that they provide a useful framework for the analysis of dynamic processes. There are many different types of network models – some theoretical, some computational – that vary with regard to a wide range of parameters, but most network models are used to explain some kind of development.

The grammar network I propose is inspired, to some extent, by a family of network models known as neural networks, or connectionist models (Rumelhart and McClelland 1986a; Elman et al. 1996). Neural networks are computational models that are widely used by cognitive scientists for modeling cognitive processes, including language use and language acquisition (e.g., Rumelhart and McClelland 1986b; Elman 1990; Chang et al. 2006; Christiansen and MacDonald 2009). One feature that makes neural networks interesting for the usage-based analysis of language is that the links between nodes have "weights," or activation values, that are shaped by processing. In general terms, the more often a particular link, or a particular pattern of links, is processed, the stronger are the weight(s) of the connections and the higher is the probability that these connections (or links) will be reused in the future. This is one way in which many network models can change. In addition, network models may change by means of particular mechanisms that create, or delete, novel nodes and novel connections or that serve to reconfigure an existing constellation of nodes and connections (Buchanan 2002).

Network models provide a very flexible tool to explain development but will only be useful if they are theoretically motivated and constrained so as to generate particular hypotheses that can be tested. In the current case, the network architecture of grammar is motivated by the cognitive organization of grammatical categories and constructions and the network analyses presented in this book will allow us to make specific predictions that can be tested in experiments. With these general comments in mind, let us now turn to the particular architecture of the grammar network I propose.

2.3   A Nested Network Model of Grammar

Building on research in construction grammar (e.g., Goldberg 1995; Croft 2001; Bybee 2010; Hoffmann and Trousdale 2013; Hilpert 2014), usage-based morphology (e.g., Bybee 1985; Hay 2001; Hay and Baayen 2005) and cognitive psychology (e.g., Bates and MacWhinney 1989; Elman et al. 1996; Tomasello 2003; Steels 2011; Ellis et al. 2016), the current study outlines a network model of grammar in which all grammatical concepts are defined by particular types of links, or relations, that indicate associative connections between different aspects of a speaker's linguistic knowledge. Specifically, I propose a "nested network model" in which the nodes at one level of analysis are networks at another level of analysis.

The model crucially relies on the notion of construction. As pointed out in Chapter 1, constructions serve functions similar to those of traditional grammatical rules; but while grammatical rules are commonly defined as algorithms that operate over primitive concepts, constructions are meaningful templates that include slots for other linguistic expressions (Langacker 1987; Goldberg 1995; Croft 2001).

Two basic types of constructions can be distinguished: morphological constructions, which are multimorphemic words such as *teacher*, consisting of a slot for a verb stem and the nominalizing suffix *-er*, and syntactic constructions consisting of two or more slots for (free) lexemes. The genitive construction (e.g., *John's car*), for instance, is an example of a syntactic construction (§1.2.3). Other examples of syntactic constructions include questions (*Where is she?*), relative clauses (*The man I met*), prepositional phrases (*on the table*) and resultative clauses (*John painted the door red*).

Since all of these structures combine a particular structural pattern with meaning, constructions are commonly characterized as signs, parallel to lexemes. In fact, in some of the construction-based literature, the notion of construction has been extended from its traditional use for particular grammatical patterns to simple lexemes and bound morphemes (Goldberg 1995: §1; Croft and Cruse 2004: §9; Hilpert 2014: §1). In this (extended) use, all linguistic signs are subsumed under the notion of construction, which is then commonly characterized as a continuum ranging from simple morphemes to complex syntactic patterns (Goldberg 2006: 5). However, in keeping with the traditional notion of construction, we will restrict the use of this term to grammatical patterns that involve at least two meaningful elements, e.g., two morphemes, words or phrases, and will use the term lexeme for monomorphemic words and single morphemes (see Langacker 1987: 83–87). As we will see, since lexemes and constructions are learned and processed in very different ways, it is reasonable to keep them separate. Still, while there