

Oxford Handbooks Online

Data in Construction Grammar

Stefan Th. Gries

The Oxford Handbook of Construction Grammar

Edited by Thomas Hoffmann and Graeme Trousdale

Print Publication Date: Feb 2013

Online Publication Date: Dec
2013

Subject: Linguistics, Morphology and Syntax

DOI: 10.1093/oxfordhb/9780195396683.013.0006

Abstract and Keywords

This chapter examines the types of data used in constructionist approaches and the parameters along which data types can be classified. It discusses different kinds of quantitative observational/corpus data (frequencies, probabilities, association measures) and their statistical analysis. In addition, it provides a survey of a variety of different experimental data (novel word/construction learning, priming, sorting, etc.). Finally, the chapter discusses computational-linguistic/machine-learning methods as well as new directions for the development of new data and methods in Construction Grammar.

Keywords: judgement data, observational data, experimental data, computational approaches, statistics

6.1. Introduction

Over the last approximately twenty-five to thirty years, a new family of linguistic theories has established itself as a powerful alternative to the then dominant generative approach to language in general and grammar in particular, the family of Construction Grammars. Many of these theories share most of their assumptions with the approach of Cognitive Linguistics, of which some are, in some sense, the grammatical part. However, the commonalities do not end there, and I want to point out two additional ones that bear on the data and methodology.

First, just like semantic categories are described in Cognitive Linguistics as radial categories—categories whose members may be linked not directly but via family resemblances—different instantiations of Construction Grammar constitute a radial category of the same kind. These Construction Grammars include Goldberg and Lakoff's Cognitive Construction Grammar (cf. Boas, this volume), Bergen's Embodied Construction Grammar (cf. Bergen and Chang, this volume), Croft's Radical Construction Grammar (cf. Croft, this volume), and others.

Thus, although Construction Grammars share many assumptions, they also differ in various ways. However, and this is the second point, given the affinity to Cognitive Linguistics, many (in particular many Construction Grammars such as Cognitive Construction Grammar, Embodied Construction Grammar), but not all (e.g., Berkeley Construction Grammar or Sign-Based Construction Grammar), explicitly commit to what Lakoff (1990: 40) has referred to as the cognitive commitment, namely “a commitment to providing a characterization of general principles for language that accords with what is known about the mind and brain from other disciplines.” Given this commitment, it is not surprising that, empirically speaking, Construction Grammar as a family of closely related grammars is probably one of the methodologically most pluralistic fields, as it utilizes a large number of different data and methodologies.

A very widespread classification of types of linguistic data has been to distinguish introspective, observational, and experimental data, where introspective data result from a speaker's second-level attention or second-level consciousness (to use Talmy's 2007 terminology), where observational data result from recordings or corpora of linguistic production in (often) noisy naturalistic settings, and where experimental data result from subjects'

behavior in designed controlled experimental situations facing carefully developed experimental stimuli. However, even though this classification is widely used, it is also a bit too simplistic, since a variety of data-gathering types exhibits characteristics of more than one of these groups. In an attempt to develop a more fine-grained classification, Gilquin and Gries (2009) develop a continuum of kinds of observational and experimental linguistic data. While, for the sake of simplicity, they present their continuum as one-dimensional, it actually integrates several different dimensions:

- *How natural does the subject perceive his (experimental) setting?*
 - most natural, e.g., speakers who know each other talk to each other in unprompted authentic dialog;
 - intermediately natural, e.g., a speaker describes pictures handed to him by an experimenter;
 - least natural, e.g., speaker lies in an fMRI unit undergoing a brain activity scan while having to press of one of three buttons in responses to digitally presented black-and-white pictorial stimuli.
- *What (linguistic) stimulus does the subject act on?*
 - most natural, e.g., speakers are presented with natural utterances and turns in authentic dialog;
 - intermediately natural, e.g., speakers are presented isolated words by an experimenter in an association task;
 - least natural, e.g., speakers are presented with isolated vowel phones.
- *What (linguistic) units/responses does the subject produce?*
 - most natural, e.g., subjects produce natural and unconstrained responses to questions;
 - intermediately natural, e.g., speakers respond with isolated words (e.g., to a definition);
 - least natural, e.g., speakers respond with a phone out of context.

For example, a situation in which a subject sits in front of a computer screen with an eye tracker, is auditorily presented with a word, is visually presented with two pictures (one of which represents an instance of what the auditorily presented word means), and responds by saying “yes” or “no” to the question whether he sees an instance of what the word refers to, can be classified as

- a fairly unnatural experimental *setting*: sitting in front of an eye-tracker;
- an intermediately natural *stimulus*: isolated words and picture input;
- an intermediately natural *response*: an isolated “yes” or “no”.

As another example, consider the case of analyzing a *there*-construction in corpus data from authentic conversations, which could be classified as

- a very natural (experimental) *setting*: an authentic conversation;
- a very natural *stimulus*: the previous turn;
- a very natural *response*: a speaker's (response) turn.

While the above classification by means of the three dimensions is neither completely exhaustive nor uncontroversial, it allows for a heuristically valuable classification of most empirical approaches in Construction Grammar in particular and probably in linguistics in general.¹ As will be shown, linguists working in Construction Grammar have used data and methods from all six extremes of these three dimensions. In fact, the amount of rigorous empirical work in the field is very high, as is indicated by the fact that, for instance, the vast majority of studies in *Constructions* and in the first volume of the new journal *Constructions and Frames* use corpus data.

In the following sections, I will discuss a range of data and methods with an eye to exemplifying how different methods have given rise to different data, and how these have advanced different subfields, or areas of application, of Construction Grammar. The structure of the exposition below is as follows: Section 6.2 will very briefly discuss a few classic studies that are based on introspective data but that have still helped found and develop Construction Grammar. Section 6.3 will discuss data and methods traditionally referred to as observational,

that is, data and methods that score highly on all three naturalness dimensions underlying the above continuum of linguistic data. Since these corpus-based approaches do not differ much with regard to these dimensions, I will instead divide them according to the ways in which the frequency data from corpora are used. Section 6.4 will then deal with what are traditionally called experimental methods, where I will distinguish different experimental approaches on the basis of the above three dimensions. Section 6.5 will very briefly address computational-linguistic/machine-learning types of approaches. Section 6.6 will conclude and present a few directions for further evolution and maturation of data and methods in Construction Grammar.

6.2. From Introspective Judgments to Other Data

It is probably fair to date the emergence of Construction Grammar to the late 1980s, when proponents of what are now regarded as different members of the family of Construction Grammars published highly influential studies. For what is now often referred to as Cognitive, or Goldbergian, Construction Grammar (cf. Boas, this volume), Lakoff's (1987) study of *there*-constructions broke important ground in the way it showed how different *there*-constructions (types of deictic and existential constructions) form a radial category of the same type as the senses of polysemous words form radial categories. Similarly, for what is now often referred to as (Berkeley) Construction Grammar, Fillmore, Kay, and O'Connor's (1988; Fillmore, this volume) study of *let alone* paved the way for many important later studies. Crucially, these and many other ground-breaking studies were ultimately based on introspective judgments about what one can and cannot say in different circumstances or contexts, with very rare references to authentic/natural examples (e.g., Fillmore, Kay, and O'Connor 1988: 524). The same is true of other earlier influential publications. A quick glance at the early volumes of *Cognitive Linguistics* indicate, for example, that the first explicitly construction-based studies (e.g., Goldberg 1992; Smith 1994; Kemmer and Verhagen 1994; Dancygier and Sweetser 1997; Morgan 1997) were all introspection-based and, if they used the word *data*, they used it referring to introspective judgments and/or example sentences (just like Lakoff 1987).

While this does in no way diminish the way these and many other publications from that time gave rise to a new field within Cognitive Linguistics, it did nevertheless not take long for many scholars to also use more and more diverse methods. In 1998, Croft and Sandra debated the degree to which linguists and their methods can contribute to matters of mental representation, with Croft arguing that "evidence from actual usage or psycholinguistic experiments" is needed to address mental representation, and 1998/1999 saw the first construction-based publications in *Cognitive Linguistics* that used data from experiments, or corpus data (or data from both): Tomasello and Brooks's (1998) experimental study of the early acquisition of Transitive and Intransitive constructions (whose discussion also involves Tomasello's earlier diary-based data), Palancar's (1999) comprehension experiment involving hitting constructions, and Gries's (1999) corpus analysis and acceptability rating experiment of verb-particle constructions. The following two sections will therefore be concerned with observational and experimental approaches that have been used in Construction Grammar studies.

6.3. Observational Approaches

As mentioned above, observational data in the form of corpus data have been playing a very important role in Construction Grammar for many years now. These corpus data do not differ much in terms of the above three dimensions of linguistic data, and corpus-based studies in Construction Grammar have used many different kinds of corpora or textual databases (space only permits mention of maximally two authors):

- in terms of *languages*: Czech (cf. Fried 2009b), Danish (cf. Hilpert 2008), Dutch (cf. Coleman 2009a), English (cf. Gries 2003a, b; T. Hoffmann 2006), Finnish (cf. Kolehmainen and Larjavaara 2004), French (cf. Chenu and Jisa 2006; Marandin 2006), German (cf. Diwald 2006; Hilpert 2009), German vs. English (Boas 2003), Greek (cf. Katis and Stampouliadou 2009), Hindi (Budwig, Narasimhan, and Srivastava 2006), Mandarin (cf. Chen 2006), Russian (cf. Eckhoff 2006), Spanish (cf. González-García 2006), and Swedish (cf. Hilpert 2008);
- in terms of *modes/registers*: journalese (cf. Croft 2009c), internet data (Stefanowitsch 2011b), but probably most studies are based on a mixture of spoken and written data that characterizes most contemporary corpora;
- in terms of *dialects/varieties*: Asian Englishes (cf. Mukherjee and Gries 2009), Belgian vs. Netherlandic Dutch (cf. Grondelaers, Speelman, and Geeraerts 2007), British vs. American English (Gries and Stefanowitsch 2010), Lancashire dialect (cf. Hollmann and Siewierska 2007),...

– in terms of *synchronic* corpora (cf. Gries 2003b) vs. *language acquisition* corpora (cf. Goldberg 1999; Diessel and Tomasello 2005b) vs. *diachronic/historical* corpora (cf. Hilpert 2008; Fried 2009b),...

While most corpus studies' data are from the most natural end on each dimension, such studies exhibit considerable variation in terms of how the corpus data are used, and since corpus data provide nothing but frequencies of (co-)occurrence, these corpus studies can be located on a (nonevaluative) cline of statistical complexity. The next few subsections will discuss differently quantitative corpus-based approaches within different areas of Construction Grammar.

6.3.1 Frequencies of (Co-)occurrence

The simplest approach to include corpus frequencies involves merely checking whether a particular construction, a combination of constructions, or a particular lexical item in a construction is attested or not. In other words, the relevant frequency distinction is between zero and one or more. While this approach may not seem particularly exciting, it can have important implications. One kind of such implications has to do with the fact that corpora may provide counterexamples for hitherto widely accepted claims. For example, the probably most widely studied Argument Structure construction, the English Ditransitive construction V NP₁ NP₂, has often been claimed to exhibit some puzzling lexical (dis)preferences, for example, that it does not occur with the verb *donate* in the ditransitive slot. However, this assessment has nearly exclusively been based on linguists' armchair judgments. Stefanowitsch (2011b), on the other hand, shows that not only do websites with *uk* as their top-level domain name contain a "substantial number" of ditransitives with *donate*, these matches also exhibit a noteworthy semantic patterning that fits into a more general account of the dative 'alternation' and its information structure. See Stefanowitsch (2011b) for more examples involving other verbs in the ditransitive.

In spite of their statistical simplicity, raw frequencies can also be highly revealing in first-language (L1) acquisition and second/foreign-language (L2/FLA) learning contexts. Regarding the former, L1 acquisition contexts, Goldberg (1999) discusses frequencies of different verbs from L1 acquisition corpora from the CHILDES database with an eye to how the high frequencies of particular (often semantically light) verbs facilitates the acquisition of argument structure constructions whose meanings are compatible with some of these verbs. Observing essentially Zipfian distributions of verbs in particular constructions, she argues, for example, that the high frequencies of *go*, *put*, and *give* facilitate the acquisition of the Intransitive-motion, the Caused Motion, and the Ditransitive constructions, respectively (cf. also Goldberg, Casenhiser, and Sethuraman 2004; Goldberg 2006a); cf. Tomasello (2003) for a book-length treatment with many different insightful case studies. Regarding the latter, L2/FLA contexts, similar observations were made in Ellis and Ferreira-Junior's (2009a) study of longitudinal data from the ESL data of the European Science Foundation (ESF) corpus with regard to *put* in Caused Motion constructions and *give* in Ditransitive constructions (cf. also Ellis and Ferreira-Junior 2009b).

6.3.2 Conditional Probabilities (Unidirectional)

The next step on a cline of statistical complexity leads to approaches involving the computation of conditional probabilities or other unidirectional measures that are based on them.² Again, these are widely used in studies in L1 acquisition and L2/FL learning, and I will again use examples from two prominent figures in these fields. As for the former, Goldberg, Casenhiser, and Sethuraman (2004) addresses the question of how reliable constructions are as predictors of sentence meaning (for an experimental approach toward that question, cf. below). They counted all instances of caused-motion meanings in the Bates corpus from the CHILDES database and then computed the cue validity of the pattern V-Obj-Loc for caused-motion meanings *p* ('caused-motion' | V-Obj-Loc). They find a high cue validity (between 0.63 and 0.83, depending on how inclusive a definition of the meaning of caused-motion is adopted), which shows that the V-Obj-Loc pattern is a good cue for the meaning it is associated with in Construction Grammar accounts. However, the more crucial implication of this finding only arises when the pattern's cue validity for 'caused-motion' is compared with the (weighted) cue validity of verbs for the same meaning, 0.68. Hence, using corpus data, Goldberg, Casenhiser, and Sethuraman (2004) showed that syntactic patterns are just as reliable as cues to sentential meaning as verbs.

As for the latter, Ellis and Ferreira-Junior (2009b) study the effect of type/token frequencies of words in slots of the Intransitive Motion, the Caused Motion construction, and the Ditransitive construction in L2/FLA in the ESF corpus.

To quantify the unidirectional association of the words to the constructions, they use—unlike most such studies, which used a bidirectional collostructional analysis measure (cf. the following section)—a very interesting unidirectional measure called ΔP , which is computed on the basis of conditional probabilities. More specifically, ΔP is the difference of the probability of an outcome O given a cue C minus the probability of O given the absence of C . They find that the first-learned types in each slot of each of the constructions—esp. the verbs—are highly distinctive for their constructional slots (both in terms of ΔP and the collostructional measures discussed below). This finding in turn supports an understanding of constructional acquisition as dependent on a larger variety of factors than are often discussed: while type and token frequencies do play important roles, the distributions of frequencies, as well as the distinctiveness of elements for the positions in which they are used and the degree to which they form chunks are also highly relevant. (Cf. Bybee and Scheibman 1999 and Bybee and Thompson 2000 on how high (type/token) frequencies of words in particular contexts/with particular meanings are related to chunking, grammaticalization, and phonological reduction processes.)

6.3.3 Association Strengths (Bidirectional)

One of the most widespread corpus-based methodological approaches in Construction Grammar is referred to as collostructional analysis, a family of several different methods. Since this approach is dealt with in a separate chapter (cf. Stefanowitsch, this volume), I will not discuss it in great detail, but in order for this chapter to be sufficiently self-contained, a few, more general remarks about this family of methods are in order.

Just like Ellis and Ferreira-Junior's (2009b) measure of ΔP , the (earlier) approach of collostructional analysis is a way to quantify association strength, which is ultimately based on collocational approaches from corpus linguistics. But unlike ΔP , it is an approach to compute a bidirectional association measure. (It is worth pointing out though that collostructional analysis does not *require* bidirectional measures; from that point of view, Ellis and Ferreira-Junior's approach is a particular implementation of collostructional analysis.) Three different methods are distinguished, most of which can be computed with Coll.analysis 3.2a (Gries 2007), a script available from the author's website:

- *collexeme analysis*, which computes for n words how strongly these words are attracted to a slot in a construction (cf. Stefanowitsch and Gries 2003);
- (multiple) *distinctive collexeme analysis*, which computes for n words how strongly these words are attracted to two or more functionally similar constructions (cf. Gries and Stefanowitsch 2004a);
- *co-varying collexeme analysis* (item-based and system-based), which computes for n words in one slot of a construction how strongly these words are attracted to the y words in another slot of the same construction (cf. Gries and Stefanowitsch 2004b; Stefanowitsch and Gries 2005).

All of these methods provide rankings of how much words and particular slots of constructions attract each other (viz., the bidirectional nature of these association measures) and what that reveals about constructional semantics, and they have given rise to many studies: Gilquin (2006) on English periphrastic causatives, Wulff (2006, 2008a) on *go (and) V* and *go/come/try (and) V*, Hilpert (2006a) on the diachronic development of verbal complements of *shall*, Hilpert (2008) on the diachronic development of Future constructions in Germanic languages, and many more. This approach also has some psycholinguistic relevance, since the preference of verbs to occur in particular (Argument Structure) constructions (i.e., the verbs' subcategorization preferences) are known to strongly correlate with linguistic processing (cf. Garnsey et al. 1997; Stallings, MacDonald, and O'seaghdha 1998; Hare, McRae, and Elman 2003; Melinger and Dobel 2005). While this must suffice here for a discussion of collostructional analysis, the topic will be revisited in the next section to discuss experiments that tried to validate this approach experimentally.

6.3.4 Multifactorial and Multivariate Approaches

Given (1) the obviously multifaceted nature of language and its relation to, or interaction with, cognitive processing and (2) the complexity and noisiness of data obtained from corpora, it is often necessary to resort to statistical methods that can do better justice to the observed facts. In theory, of course, nearly every phenomenon studied corpus-linguistically can, and probably should, be studied multifactorially, so the range of possibilities that could be surveyed is extremely large. I will mention two kinds of approaches, which are not only multifactorial/-variate but also methodologically pluralistic in how they combine data from corpora and data from experiments.

The first of these is concerned with a notion from the very beginning of Construction Grammar, idiomaticity. As mentioned above, early studies in Construction Grammar were devoted to the study of different kinds of idioms and to how the study of these items that are often considered ‘marginal’ illuminates the study of more regular constructions. However, as has been well known, idiomaticity is a perplexingly multidimensional notion, hard to operationalize or even just rank-order on the basis of introspection alone. Wulff (2009) is a study that approaches idiomaticity on the basis of experimental and corpus data for thirty-nine V-NP idioms from the British National Corpus. First, she collected idiomaticity judgment data from subjects using the method of magnitude estimation. Second, and more importantly here, she used two different corpus-based ways to operationalize different dimensions of idiomaticity: collocational overlap and a measure of formal flexibility that was in turn based on twenty idiomatic variation parameters (describing morphological and syntactic parameters of the idioms’ use). Using multivariate and multifactorial methods—principal components analysis and multiple regression—she then identified which idiom variation parameters cluster (and are thus likely to underlie perceptions of idiomaticity) and validated these factors/clusters on the basis of the speaker judgments. This showed, among other things, that compositionality was not as strong a predictor as was commonly held; cf. Wulff (2008b) for more detailed discussion.

The second multidimensional study to be discussed here was concerned with identifying prototypical instances of constructions. Gries (2003b) retrieved examples of the dative alternation from the British National Corpus and coded them for a large number of morphological, syntactic, semantic, and discourse-pragmatic characteristics. He then used a linear discriminant analysis to determine which of these characteristics, if any, were good predictors of the constructional choices in the corpus data. He showed that nearly 89% of all constructional choices could be classified correctly (and how important each variable’s contribution to that was), but more importantly, each corpus instance was assigned a discriminant score that reveals how good, or prototypical, an example of the ditransitive and the prepositional dative is (in terms of how sure *and* correct the analysis was in assigning a constructional choice). As a first attempt at validation of these corpus-based findings, he discussed several salient constructions—prototypical examples that were predicted correctly and examples where the model was wrong—but the more forceful validation was an acceptability judgment experiment, in which speakers rated sentences well when they occurred in the construction that the corpus-based discriminant analysis predicted for them. Among other things, Gries argued, therefore, that this kind of corpus-based multifactorial approach is a valid and useful tool to obtain goodness-example information for data that can be useful for, say, acquisition approaches or the study of alternations, etc. For other multifactorial corpus-based applications, cf. Gries (2003a), Brenier and Michaelis (2005), Bresnan et al. (2007), or T. Hoffmann (2006, 2011).

6.4. Experimental Approaches

Apart from a large and growing number of corpus-based approaches, studies in Construction Grammar have also employed many different kinds of experiments. While these are typically not from the most technical/artificial type of settings, they nevertheless exhibit quite some variation. This section discusses several experimental approaches with an eye to surveying the kinds of methods and data that were used, and it does so by moving roughly from more natural/less artificial settings, stimuli, and responses to less natural/more artificial ones.

The experimental approaches that are among the most natural on all dimensions are those involving young children. For example, in one of the first nonintrospective Construction Grammar papers in *Cognitive Linguistics*, Tomasello and Brooks (1998) performed an experiment in which children (mean age 2;3) were taught novel verbs, one with/in a transitive, the other with/in an intransitive meaning/scenario. They then encouraged the children’s use of the novel verbs in a construction by, for example, saying “This is called *meeking*. Can you say *meeking*? Say *meeking*?” and by asking descriptive questions in an elicitation task. While these tasks would be somewhat artificial in an adult-only context, interactions like these are, of course, not at all rare in discourses between children and their caretakers, which is why this kind of experimental approach is categorized as natural on all dimensions. Tomasello and Brooks find that “young children learn their first sentence-level constructions on a verb-specific basis” (1998: 391), supporting the item-based approach for which the Tomasello and his research group have become so well-known (cf. Abbot-Smith and Tomasello 2006, 2010; Kidd, Lieven, and Tomasello 2010, to name but a few more examples).

A similar example is Casenhiser and Goldberg’s (2005) study of novel verb learning. Just like Tomasello and Brooks,

they exposed (somewhat older) children (mean age: 6;4) to nonce verbs, but, unlike Tomasello and Brooks, they exposed them to nonce verbs in a phrasal pattern that does not exist in English *and* they systematically varied the token frequencies with which the nonce verbs occurred in the novel pattern. After the short training session (less than three minutes), the children participated in a forced-choice comprehension task; the dependent variable was whether they could understand sentences using the novel pattern correctly, especially when their training involved a token frequency distribution that was skewed in a way that is skewed similarly to the Zipfian distributions of verbs in constructions discussed above, showing that children are very fast at identifying probabilistic patterns in skewed distributions and associating a meaning with them.

Another range of experimental approaches used in Construction Grammar involves several paper-and-pencil tasks, which involve intermediately unnatural experimental settings but that differ with regard to the naturalness of the stimuli and the 'output' produced by the subjects.

One set of experiments that has provided different kinds of useful findings involves priming effects. Many priming studies have restricted themselves to a purely syntactic/structural view of priming, but in an important study Hare and Goldberg (1999) extended a previous study by Bock and Loebell to determine to what degree, if any, priming may not just be syntactic but also influenced, or reinforced, by semantic factors. Hare and Goldberg used a picture-description task in which subjects describe ditransitive scenarios after having heard one of three different prime sentences (or an intransitive control sentence). Importantly, they found that "the order of expression of coarse semantic roles" and "the level of the mapping between semantics and syntax" influenced subjects' reaction (cf. also F. Chang, Bock, and Goldberg 2003).

Other priming studies involve, for example, experimental designs where the subjects do not produce a full 'normal' sentence but a slightly less natural response, namely where they complete a sentence fragment. One such example involves foreign language learners' knowledge of constructions. Gries and Wulff (2005) conducted a sentence-completion experiment with advanced German learners of English. In this study, primes were set up to bias subjects into producing either ditransitives or prepositional datives to determine whether (1) German learners exhibit the same kinds of priming effects as native speakers and (2), just as importantly, whether German learners exhibit the same kinds of verb-construction preferences as native speakers of English. Gries and Wulff found both of these effects: the learners exhibited constructional priming effects **and** verb-construction preferences that were very similar to native speakers, but they also showed that the verb-construction preferences they found were not due to translational equivalents' transfer effects. Gries and Wulff (2009) then conducted a similar study, this time testing for whether priming can be obtained for two complementation patterns—*to-* vs. *ing-*complementation after verbs such as *like* or *try*—and, if so, what the source of the priming effect is. They again found strong and significant priming effects for both constructions, mainly from the verb in the prime but also the subjects' own last completion. Both studies therefore lent support to exemplar-based approaches toward linguistic knowledge in general and collostructional knowledge/subcategorization preferences in general.

Another example is Gries, Hampe, and Schönefeld's (2005) study of *as*-predicatives. They were concerned with the question of what kinds of frequency data are most useful to the study of verbs in constructions. They first undertook a corpus study of the *as*-predicative (as in, e.g., *He regarded that as a big mistake*) to determine verbs that are frequent or not so frequent in that construction, as well as verbs that are highly attracted or barely attracted to the construction (in terms of collostructional attraction). Then, they presented subjects with sentence fragments featuring verbs from each of the four groups that resulted from crossing the frequency and the attraction conditions. The dependent variable was therefore whether subjects would use an *as*-predicative or not, and they found that the collostructional measure had a very large effect on the subjects' completion patterns (as had the voice of the sentence fragment) whereas raw frequency did not, which lends experimental support for corpus studies of constructions using uni- or bidirectional measures of association.

Studies in which subjects were requested to do something less natural than produce or complete sentences, include cases where subjects fill gaps or sort sentences. As for the former, Dąbrowska (2009) presented native speaker subjects with sentences from dictionary definitions of verbs of walking from which these verbs have been omitted. She asks them to fill the gap, a not particularly natural response type, and finds that subjects are quite good at finding the right verb on the basis of the collocational knowledge they have accumulated over time.

As for the latter, Bencini and Goldberg (2000) used a sorting paradigm to study which components of a sentence—

the main verb or the Argument Structure construction—are most central to the sentence's overall meaning. Native speakers of English received sixteen cards, each with a different sentence that used one of four verbs in one of four Argument Structure constructions (Ditransitive, Transitive, Caused Motion, and Resultative construction); the stimuli can therefore be categorized as rather natural. The subjects were then asked to sort the sixteen sentences into piles depending on overall similarity of meaning (i.e., perform a not-so-natural linguistic task). The dependent variable and the question in point was whether the subjects would produce piles based on the verbs or on the constructions. It turned out that the subjects produced significantly stronger construction-based clusters, which underscored the relevance of Argument Structure constructions for sentence meaning.

A replication of this study provided additional results. Gries and Wulff (2005) replicated this experiment with advanced German learners of English, with additional findings. The German learners also exhibited a significant preference for construction-based sortings—in fact an even stronger effect in this direction than the native speakers, but Gries and Wulff also analyzed the sorting data by means of exploratory data analysis methods, a hierarchical agglomerative cluster analysis and a principal components analysis. Both of these methods not only supported the findings that the sorted sentences came in construction-based piles (accounting for more than 90% of the variance in the data) but also resulted in a dendrogram that reflects how similar the constructions are to each other in the eyes of the subjects. Interestingly enough, the clustering of the constructions is perfectly compatible with their theoretical treatment in Construction Grammar such that, for example, Resultative and Caused Motion constructions are related most strongly, reflecting Goldberg's (1995) analysis. This is therefore a case where a more comprehensive statistical analysis of the data could yield results that even go beyond the original question.

An experimental approach that is similar in terms of experimental setting and input, but involves the subjects' producing maybe more unnatural responses are experiments involving judgment data. Dąbrowska (2008) studied questions with long-distance dependencies and finds, using acceptability judgments, that these questions exhibit very strong prototype effects such that questions that correspond closely to one of several corpus-derived templates receive significantly better ratings. As mentioned above, Gries and Wulff (2005) determined that their learners of English had the same verb-construction preferences as native speakers by comparing their experimental behavior to the verbs' and their translational equivalents' preferences in English and German corpora. Gries and Wulff (2009) tested the German learners' preferences more directly. Subjects were presented with sentences of the two complementation patterns V_1 to V_2 and V_1 V_2 -*ing*, but the sentences were designed to contain V_1 's that collostructionally preferred the first or the second pattern in corpus data. The dependent variable was the subjects' acceptability ratings to the four combinations (of two constructional preferences and two constructional stimuli). They found again that the learners were very sensitive to the verbs' constructional preference, giving high ratings to stimuli where verbs were used in the construction they prefer, and low ratings otherwise, also lending support to the kind of assumption made in exemplar/usage-based models that even learners are able to keep track of the frequencies with which verbs are used in construction.

The final experimental method to be discussed here involves a design with a rather artificial design (on all three levels). Gries, Hampe, and Schönefeld (2010) conducted a follow-up study of their 2005 sentence-completion experiment, which involved a self-paced reading task. On the basis of a larger corpus sample, they again crossed frequency of co-occurrence (high vs. low), collostructional attraction (high vs. low), and voice and presented subjects with sentences from the British National Corpus that contained these verbs but were altered to render their lengths and complexities comparable, as well as replacing context-dependent expressions such as proper names by more generic expressions. The subjects read the sentences word-by-word such that they had to press a button to request and obtain the next word. The dependent variable was the time from the presentation of one word till the request of the next word. With only few subjects, they obtained 254 reading times, but when they analyzed the reading time of the word following—the word that should reveal to the subjects whether their initial parse expectation based on the verb was correct or not—they found that again frequency had no significant effect at all ($p = 0.293$), whereas collostructional attraction exhibited a marginally significant effect in the predicted direction ($p_{\text{two-tailed}} = 0.065$), again supporting the importance of association strengths over raw frequencies.

Space does not permit discussion of more experimental paradigms that would do merit to their complexity and potential, but a final group of experiments must nevertheless not go completely unmentioned, namely the large body of work that has been done in the areas of Simulation Semantics and Embodied Construction Grammar (cf. Bergen and Chang, this volume). Consider as a case representative for much work in these fields a very interesting

study by Bergen and Wheeler. Starting out from the view that understanding language often involves mental perceptual and motor simulations (as indicated by activation of areas in the brain responsible for motor action), they test action-sentence compatibility effects, that is, whether the direction of motion represented in a sentence is compatible with the hand movement the subjects have to make to press a response button (and thus speeds up reaction times) or not (and thus slows down reaction times). Using this paradigm, they find that progressive aspect and perfective aspect result in very different Action-sentence compatibility effects, which suggests (1) that the different aspects result in different mental simulations of the actions described and (2) that grammatical features such as aspect modulate “second-order properties of the mental simulation to be performed” and “what part of an evoked simulation an understander focuses on, or the grain of detail with which the simulation is performed” (Bergen and Wheeler 2010: 155). Studies like this are still rather rare but point to very intriguing possibilities for future research along these lines; cf. Bergen (2007) for an excellent summary of different experimental paradigms in these areas, which are evolving quickly and becoming more and more relevant to the field of Construction Grammar.

6.5. Computational-Linguistic/Machine-Learning Approaches

The kind of data and methodology that are least used in construction-based approaches are computational-linguistic approaches involving, for example, machine-learning or simulation-based approaches, and much of the work in these areas that would in fact be relevant to construction-based approaches does not establish a direct connection to Construction Grammar. One example is F. Chang et al. (2000), who developed a connectionist model to test whether structural/syntactic priming—which, as discussed above, is seen by some as constructional priming—can be considered as resulting from implicit learning (rather than, say, from residual activation of nodes in a spreading activation network). They find that their type of simple recurrent network suggests that priming may indeed result from the very same mechanisms that underlie language learning in the first place, and that, among other things, a model that involves/simulates message comprehension yielded more priming effects of the type that humans exhibit.

An approach that is less computational and, thus, more transparent to the traditional linguist is the Traceback approach developed by Dąbrowska, Lieven, and colleagues (cf., e.g., Dąbrowska and Lieven 2005; Lieven et al. 2009; Vogt and Lieven 2010 for recent examples). In this approach, a program called Autotracer identifies all multiword utterance types in a test corpus, typically the last two hours of recorded speech of a child, and then identifies all (continuous and discontinuous) strings that occur at least twice in the prior recordings and that contained overlapping lexical material with the target utterance types in the test corpus. After all potential component units were identified, the program attempts to build up all utterance types in the test corpus from the potential component units with superimpositions, substitutions, and additions. The objective is to determine how many of the novel utterances of a child can actually be traced back to only slightly changed previous utterances, and findings indicate that, in spite of the sparsity of even the densest language acquisition corpora, often the vast majority of children's novel utterances can be accounted for as exact repetitions or with one operation. In addition, results suggest that children are in fact learning chunks—and do not freely assemble utterances from parts.

Methods such as these, or those discussed in Dominey (2006), are not yet particularly frequent in Construction Grammar, but they can be extremely useful additional tools, since they allow the researcher to identify patterns in use, as well as developmental trends in acquisition and learning, that are virtually impossible to detect otherwise.

6.6. Future Developments

As the previous sections have illustrated, Construction Grammar is an empirically and methodologically vibrant field, using different data and different cutting-edge techniques, which is probably in no small part due to the closeness of the field's ‘superfield,’ Cognitive Linguistics, to Cognitive Science, a field in which discussions of data and methods have been prominent even during the time during which much of theoretical linguistics was dominated by introspective data. However, this generally fortunate state does not mean that there are no ways in which Construction Grammar can evolve further in terms of data and methods, and the following brief discussion mentions a few directions in which the field is most likely to move.

With regard to observational approaches, recent developments in Corpus Linguistics have resulted in many

methodological innovations and techniques that are relevant to Construction Grammar in particular and Cognitive Linguistics in general. Corpus linguists are exploring:

- more and more diverse association measures to quantify if and how much different linguistic elements are attracted to each other (cf. section 6.3.3 above);
- ways to identify uninterrupted and interrupted *n*-grams, which can inform language acquisition research on multiword units (cf. section 6.5 above) but also the probabilistic identification of parts of speech (cf. Redington, Chater, and Finch 1998; Mintz, Newport, and Bever 2002);
- ways of quantifying the dispersion/distribution of linguistic elements, which can help explain the learnability of these elements (cf. Ambridge et al. 2006); etc.

Obviously, Construction Grammarians have much to gain from staying up-to-date with regard to such developments. Similar comments apply to psycholinguistic and broader cognitive-linguistic fields, where experimental methods are continuously developed and/or refined, and, for example, the collection of papers in Gonzalez-Marquez et al. (2007) discusses many potentially interesting applications such as eye movement research experiments on language and space, most of which should be applicable and useful in Construction Grammar contexts, too. For example, it is probably only a matter of (little) time until the kind of imaging techniques discussed by Coulson (2007) will be applied to more specifically construction-based questions (cf. Pulvermüller, this volume).

Finally, with the importance that usage plays in most contemporary incarnations of Construction Grammar, computational simulations of first-language acquisition or diachronic change will assume a more central role than they have done so far, and Edelman (2007) surveys some notions relevant in this context.

A final development relevant to empirical Construction Grammarians transcends the (somewhat tenuous) distinction between observational and experimental approaches: how data are analyzed statistically. There are still many studies in which data are not analyzed with the necessary degree of rigor. In addition, there are relatively new developments in the field of statistics that are very promising. One of these is the method of mixed-effects models, or multilevel models, a family of approaches of generalized linear models that is extremely powerful in how it handles random effects (such as subject-/stimulus-specific variation), unequal cell frequencies, and missing data and in how these advantages make statistical estimates much more precise (cf. Gelman and Hill 2008). While standards are still emerging in this domain, this is a methodological trend that Construction Grammarians should be and remain aware of.

To conclude, true to the spirit of the cognitive commitment mentioned above, researchers working in/on Construction Grammar already make use of a vast array of data and methods that have proven useful and yielded very informative results in many neighboring disciplines. It seems that, over time, the trend toward methods that are more rigorous and replicable than introspective judgments has only become stronger, and it remains to be hoped that the above desiderata and the adoption of some of the more recent developments in Cognitive Linguistics, Psycholinguistics, and Corpus Linguistics also find their way into the Construction Grammarian's toolbox.

Notes:

(1.) Even this lengthier characterization is still a simplification, since subjects may be presented with different kinds of stimuli at the same time, etc.

(2.) Conditional probabilities are written as $p(E|F)$, which means 'the probability of an event E, given that another event F has occurred.' An example would be the probability $p(\text{ditransitive}|\text{recipient} = \text{animate})$, i.e., the probability that a speaker will use a Ditransitive construction (as opposed to a prepositional dative with *to*) when the recipient is animate (as opposed to inanimate).

Stefan Th. Gries

Stefan Th. Gries is Professor of Linguistics at the University of California, Santa Barbara, as well as Honorary Liebig Professor at the Justus-Liebig-Universität Giessen. Theoretically, he is a cognitively and Construction Grammar oriented linguist in the wider sense of seeking explanations in terms of cognitive processes; methodologically, he is a quantitative corpus linguist at the intersection of corpus, cognitive, computational, and quantitative linguistics. He is founding editor-in-chief of the international peer-reviewed journal

Corpus Linguistics and Linguistic Theory, associate editor of Cognitive Linguistics, and on the editorial boards of CogniTextes, Constructions and Frames, Corpora, Language and Cognition, and others.