Frequency: Psychological and methodological considerations

Martin Hilpert (University of Neuchâtel)

Abstract

This chapter discusses the role of frequency for Construction Grammar, especially concerning usage-based models of language. For this purpose, the chapter offers definitions of different aspects of frequency, and it discusses different ways in which frequency manifests itself in the use of constructions. The chapter differentiates between token frequency, type frequency, relative frequency, frequency of co-occurrence, and dispersion. It discusses how these aspects of frequency can be measured on the basis of corpus data, and how these measurements allow the observation of different frequency effects that relate to phenomena such as entrenchment, ease of processing, productivity, phonological reduction, and resistance to regularization. These frequency effects are illustrated on the basis of experimental and corpus-based analyses of lexical, morphological, and syntactic constructions. The chapter also addresses frontiers and open questions regarding the role of frequency in constructionist research. Not only is the relation between corpus frequencies and theoretical notions such as entrenchment far from trivial, it is also important not to attribute effects to token frequency that can in fact be explained by other, correlating variables. The chapter will examine a number of strategies that can be used to go beyond the use of frequency values in the future development of Construction Grammar.

Key words: token frequency, type frequency, relative frequency, frequency of co-occurrence, dispersion, entrenchment, priming, statistical preemption, productivity

## 1 Introduction

This chapter discusses the role of frequency for Construction Grammar, especially concerning usage-based models of language (Bybee 2010). In doing so, the chapter addresses issues of language processing and production, but also concerns that are methodological, relating to the corpus-based study of constructions through various statistical techniques (Yoon and Gries 2016). The importance of frequency measurements in

constructional research has grown over recent years, reflecting the continuing evolution of the field. This development can be illustrated by comparing three definitions of the term *construction* that have been proposed by Adele Goldberg (1995, 2005, 2019), and that have had a considerable impact in the field:

(1) C is a CONSTRUCTION iff$_{def}$ C is a form–meaning pair <$F_i$, $S_i$> such that some aspect of $F_i$ or some aspect of $S_i$ is not strictly predictable from C's component parts or from other previously established constructions. (Goldberg 1995: 4)

(2) Any linguistic pattern is recognized as a construction as long as some aspect of its form or function is not strictly predictable from its component parts or from other constructions recognized to exist. In addition, patterns are stored as constructions even if they are fully predictable as long as they occur with sufficient frequency. (Goldberg 2006: 5)

(3) [C]onstructions are understood to be emergent clusters of lossy memory traces that are aligned within our high- (hyper!) dimensional conceptual space on the basis of shared form, function, and contextual dimensions. (Goldberg 2019: 7)

If the three definitions are compared side by side, several differences become apparent. The first definition places non-compositional meanings and idiosyncratic formal characteristics at its center. Constructions such as *the Xer the Yer* (Kay and Fillmore 1999) or the *way*-construction (Goldberg 1995) are well-known illustrations of these characteristics, which constitute a fundamental motivation for viewing constructions as conventionalized form-meaning pairings.

The second definition builds on the first one but explicitly includes frequency, which is understood in this context as token frequency, that is, the repeated occurrence of a linguistic pattern in language use. The definition claims that sufficient exposure to such a pattern will result in the mental representations of a construction, even if the pattern itself is transparent and rule-governed. To take an example, a string such as *I don't know* conforms to the canonical syntax of English negated declarative sentences, and it conveys a meaning that is fully compositional. Still, given its high frequency of use, speakers of English are likely

to store it redundantly in memory, as a construction. Besides semantic or formal idiosyncrasy, high frequency of use thus provides a second important criterion for constructionhood. The definition further suggests that there is a threshold that distinguishes entrenched constructions from patterns that occur more rarely and that are not mentally represented as such. Such a threshold presupposes a continuum of entrenchment between highly frequent, conventionalized constructions and rare patterns that represent chance co-occurrences of linguistic elements. Whereas *I don't know* is highly entrenched, sentences such as *I don't row* or *I don't sow* are not. With its emphasis on frequency, Goldberg's second definition reflects the increasingly close ties between Construction Grammar and both psycholinguistic work and corpus-based research.

The third definition further expands on frequency-related characteristics of constructions. Goldberg (2019: 7) states that this definition is more aligned with what is currently known about human memory, learning, and categorization. It is useful to unpack three aspects of the definition. First, constructions are viewed as clusters of memory traces, which means that repeated exposure to similar linguistic patterns eventually leads speakers to form generalizations. This can be illustrated with the ditransitive construction, which combines a subject, a verb, and two objects, as in *Mary handed me an envelope* or *We sent them a message*. These examples contain different words, but instantiate the same structural pattern. The ensuing generalization is characterized as 'lossy', so that the mental representation of the construction is more abstract than the actual instances of language use that form its basis. ~~At~~ the same time, the generalization is anchored in usage, so that linguistic elements that appear especially frequently represent the prototype of the construction (Bybee 2010: 79). Second, Goldberg situates constructions in a high-dimensional conceptual space. Instances of constructions vary along several dimensions of form and meaning. A speaker who is exposed to such variation will form generalizations on the basis of that experience. For example, a speaker of English would register that in ordinary language use, most instances of the ditransitive construction involve animate recipients, but that there are occasional exceptions, as in *We gave the bathroom a makeover*. Knowledge of constructions thus involves knowledge of variation, specifically including knowledge of the relative frequency of variants. A third aspect in which the definition breaks new ground concerns the role of context. Language use is sensitive to context, which means that frequency of use has to be understood relative to speech

situations. The constructions that a speaker is likely to encounter in one context may be near-absent in another, depending for example on the formality of the situation or the social distance between the communicators. An example for this is offered by Hay and Foulkes (2016) in a study of pronunciation variation in New Zealand English. In words such as *city*, speakers may realize the intervocalic /t/ either as a voiceless alveolar stop [t] or as a voiced or flapped variant [d/ɾ]. The extralinguistic context influences speaker behavior, so that when speakers verbalize events that lie in the more distant past, they are relatively more likely to use the more traditional, voiceless variant (Hay and Foulkes 2016: 322). The example shows that constructions are sensitive to subtle characteristics of the speech situation. In the light of this, what Goldberg's third definition accomplishes is thus a more realistic, but ultimately also more complex account of how frequencies of use shape speakers' mental representations of constructions.

To summarize what has been said so far, frequency has become an increasingly important notion in constructional studies. Whereas initially the focus was centered on token frequency, current research takes into account a much broader range of frequency-related measures that often pertain to variation in language use. It is the purpose of this chapter to present an overview of the measures and techniques that have been used in constructional analyses, to outline what results can be obtained through their use, and to discuss how these results inform constructional theories of language.

The remainder of this chapter is structured as follows. Section 2 will offer definitions of frequency, and it will discuss the way frequency manifests itself in the use of constructions. The section will differentiate between token frequency, type frequency, relative frequency, frequency of co-occurrence, and dispersion. It will be explained how these aspects of frequency can be measured on the basis of corpus data, and how these measurements allow the observation of different frequency effects (Bybee 2010, Pfänder and Behrens 2016) that relate to phenomena such as entrenchment, ease of processing, productivity, phonological reduction, and resistance to regularization, among others. These frequency effects will be illustrated on the basis of experimental and corpus-based analyses of lexical, morphological, and syntactic constructions. It will be discussed how the insights from these empirical studies feed back into the development of constructional theories. Throughout the section, the discussion will showcase corpus-based methods that draw on frequency as a means to analyze constructions. For each method that is featured, it will be

clarified what kind of frequency data enters the analysis, how that data is processed, and what can be learned from the results. Case studies from the corpus-based constructional literature will be used to illustrate relevant concepts.

Section 3 of the chapter will turn to frontiers and open questions regarding the role of frequency in constructionist research. Not only is the relation between corpus frequencies and theoretical notions such as entrenchment far from trivial (Blumenthal-Dramé 2017), it is also important not to attribute effects to token frequency that can in fact be explained by other, correlating variables. Gries (2022a, 2022b) suggests a number of avenues that should prompt constructional research to go beyond the use of frequency values, taking into account paradigmatic and syntagmatic variability as well as dispersion and contingency. These suggestions will be taken up in order to flesh out their implications for the future development of Construction Grammar.

## 2 Frequency measures and frequency effects

This section discusses five different ways in which frequency manifests itself in language use and language processing. Token frequency, type frequency, relative frequency, frequency of co-occurrence, and dispersion will be presented individually, focusing on their respective effects and their relevance for constructional analyses. The discussion will also point towards interrelations between these aspects of frequency.

### 2.1 Token frequency

Token frequency captures how often an element or structural pattern appears in language use. It is measured through counts of elements in linguistic corpora. For example, in the British National Corpus, the word *time* has a token frequency of more than 150,000, whereas the word *chronology* appears a mere 300 times. Token frequencies are commonly normalized, for example to instances per million words, in order to make them comparable across corpora. High token frequency is known to impact language processing in several ways (Ellis 2002). Five effects of high token frequency that have direct implications for

constructionist theories of language are known as chunking, entrenchment, reduction, conservation, and conservatism.

The first of these, chunking, refers to the phenomenon that speakers cognitively process a string of linguistic elements as a holistic unit (Bybee and Scheibmann 1999). This is the case for strings such as *by the way* or *you know*, which serve as discourse markers and thus convey meanings that go beyond the meaning of their component parts. *By the way* signals that the speaker temporarily deviates from the discourse topic (Fraser 2009). With *you know*, the speaker can invite the hearer to draw an inference (Jucker and Smith 1998). Such non-compositional meanings, which constitute an important definitional criterion of constructionhood, are directly related to holistic processing, which in turn is driven by high token frequency. Chunking is further relevant for the processing of complex syntactic constructions. For example, speakers of English holistically represent constructions such as *the Xer the Yer* (Kay and Fillmore 1999, Hoffmann 2019), so that when they hear an utterance that begins with *The more I think about it*, the first part of the construction will trigger expectations about how the utterance will be continued, thereby facilitating language processing. Another effect of chunking relates to syntactic constituency (Bybee 2010: 136). If a string of elements is frequently processed, speakers will treat it as a syntactic unit. Hilpert (2015: 348) offers the example of the string *sitting and waiting*, and argues that its syntactic behavior in questions such as *What are you sitting and waiting for?* is that of a single verbal unit. Questioning the prepositional object of the verb *waiting* is only possible because *sitting and waiting* is processed as a chunk. By comparison, a question such as *What are you walking and waiting for?* sounds unidiomatic, since the string *walking and waiting* is not processed as a chunk and does not function as a single verbal constituent.

A second important effect of high token frequency is entrenchment. Repeated exposure to a linguistic unit increases the strength with which that unit is mentally represented. Elements and patterns that are highly entrenched can be retrieved from memory more quickly and more accurately (Ellis 2002: 152). A finding that is particularly relevant for Construction Grammar has been reported by Arnon and Snider (2010), who conducted an experiment in which the participants had to indicate whether English n-grams like *don't have to worry* and *don't have to wait* are possible parts of grammatically formed sentences. Both fragments are indeed grammatically well-formed, but there is a difference in that the first one is more frequent than the other. The participants were quicker to

confirm the grammaticality of the more frequent n-gram. Crucially, this is not due to the token frequency of the component words and word pairs, which were controlled for in such a way that they are actually identical in frequency. Entrenchment can thus be shown to facilitate language processing even at the level of multi-word units.

Another common effect of high token frequency concerns language production. Linguistic units that are used frequently tend to become phonologically reduced. Bybee and Thompson (1997: 576) give the example of *BE supposed to*, which can be reduced to [spostə] in utterances such as *That's not supposed to happen*. Phonological reduction of this kind is common in cases of grammaticalization (Hopper and Traugott 2003: 69), as for instance in the case of *be going to* and its reduction to *gonna*, or *want to* and *wanna*. The phonologically reduced forms can be shown to establish themselves as independent constructions that differ in terms of meaning and structure from their respective sources (Lorenz 2013).

What the effects of chunking, entrenchment, and reduction show is that high token frequency is a powerful driver of language change. At the same time, high token frequency also has a conserving effect (Bybee and Thompson 1997: 577). Linguistic elements that are frequently used are relatively more resistant to analogical change, which is apparent for example in irregular past tense formations such as *keep* - *kept*. Whereas *keep* has retained its irregular past by virtue of its high token frequency, less frequent verbs such as *weep* or *leap* have succumbed to analogical pressure and have adopted regularized past tense forms. The conserving effect of high token frequency is also evident in syntax, for example in English expressions that maintain older forms of sentential negation (Tottie 1991, Bybee 2006). Whereas the canonical pattern for an English negated sentence involves *do*-support, as in *I don't believe that*, post-verbal negation is preserved with modal auxiliaries (*You should not believe that*) and in fixed expressions that involve lexical verbs (*make no mistake*, *it gives me no joy*). In these cases, high frequency of use motivates the maintenance of the older syntactic patterns.

A speaker who experiences a linguistic unit with high token frequency forms increasingly clear expectations about its distributional behavior. This is even true for children who are still acquiring their first language. Children are known to produce overgeneralization errors such as **bringed* for *brought* or **breaked* for *broke* (Clark 1987: 19). Brooks et al. (1999) showed that this phenomenon is sensitive to token frequency. Specifically, it was

shown that children are less open towards new combinations of words and constructions if the words in question are highly frequent. In the study, children were encouraged to use intransitive verbs in a transitive way. Existing intransitive verbs were extended more readily if they were infrequent. That is, while children were reluctant to transitivize frequent verbs such as *come* (e.g., *He came the car*), they showed a greater tendency to transitivize infrequent intransitive verbs with the same meaning, as for example *arrive* (e.g., *He arrived the car*).

With regard to Goldberg's (2006: 5) second definition of constructions, the effects discussed in this section motivate why token frequency is important to consider in constructional analyses. The effects also suggest that high token frequency tends to bring about non-predictable characteristics in constructions. An entrenched multi-word unit such as *I don't know* is syntactically canonical and semantically transparent, but in actual usage, that string will often be used in its phonologically reduced form *dunno* and with a non-compositional meaning. As a response to the question *Would you like to watch Netflix?,* the string *I dunno* would indicate that the speaker actually has a different preference. The inclusion of high token frequency as a definitional criterion of constructionhood is therefore fully compatible with a view of constructions as idiosyncratic linguistic units that need to be learned.

Token frequency is further relevant with regard to the influence of linguistic context on speakers' choices between alternative options. Brown and Rivas (2012) show this in a study of Spanish presentational constructions with *haber* 'have', as in *Hubo problemas* 'There are problems' (lit. *have*$_{SG}$ *problems*). Prescriptively, the presentational construction calls for a singular form of the verb, regardless of whether the noun phrase that is presented is in the singular or the plural. That said, in many varieties of Spanish, plural noun phrases may actually trigger plural agreement in the verb. Brown and Rivas (2012) demonstrate that this phenomenon is sensitive to token frequency, more specifically the frequency with which a plural noun appears as a grammatical subject in other clause-level constructions. A plural noun such as *estudiantes* 'students', which frequently appears as a grammatical subject, exhibits a stronger tendency to trigger plural agreement than a plural noun such as *fraternidades* 'fraternities', which is used less often with the function of a grammatical subject.

## 2.2 Type frequency

Type frequency represents the number of different variants of a construction in language use. More specifically, type frequency can be determined through the use of corpora, by exhaustively extracting the instances of a given construction, so that the different elements that can fill a slot of that construction can be counted. The concept can be illustrated with the English regular past tense construction, which is a pattern with very high type frequency (Bybee 1995). The suffix *-ed* appears with a large number of verbs, yielding forms such as *walked*, *painted*, or *opened*. By comparison, the type frequency of forming the past tense through a vowel contrast, such as *drink* - *drank*, *sing* - *sang*, or *swim* - *swam*, is much lower. Type frequencies can not only be determined for morphological constructions, but also for syntactic patterns. For example, Israel (1996) investigated the diachronic growth in type frequency of the English *way*-construction by examining the main verbs that appeared in that construction during different historical periods.

Goldberg (2006: 99) points out that high type frequency is connected to productivity, that is, the ease with which speakers produce and process new instances of a construction. A construction that accommodates a broad range of elements as slot fillers is more likely to be extended further to new elements, as opposed to a construction that occurs only with a select few. Corpus-based operationalizations of productivity often combine type counts with other measures, including token frequency and the relative prevalence of hapax legomena, that is, types that appear only once. For example, the measure of *potential productivity* is calculated by taking the hapax legomena of a construction and dividing that number by the overall token frequency of the construction (Baayen 2009: 902). If a construction is used with high token frequency, and a large ratio of its types occurs just once, this constitutes strong evidence that speakers regularly use that construction in creative ways.

Barðdal (2008: 34) argues that assessments of the productivity of a construction need to take into account the coherence between its types. Coherence can be defined both in terms of meaning and phonological similarity. Type frequency and coherence are seen as inversely correlated, so that if a construction occurs with many types, their coherence is likely to be low. Suttle and Goldberg (2011: 1239) make a similar point when they relate the productivity of constructions to type frequency, variability, and similarity. If a speaker

observes a pattern with a large number of types that are variable in terms of meaning, a new instance that is similar to any one of those types will have a good chance of being acceptable.

An effect of type frequency can be seen in a study by Brooks and Tomasello (1999), who presented children between 2 and 4 years of age with nonce verbs such as *meeking* and *tamming*. The respective verbs were presented exclusively in either the active voice or the passive voice. After the training phase, the children were encouraged to use the verbs themselves, and they were provided with contexts that would favor either the use of the active or the passive. Brooks and Tomasello (1999: 34) observe that the children are generally reluctant to generalize the verbs from one construction to the other, but they further note an asymmetry between the two constructions, in that generalizations towards the passive, which is characterized by a relatively lower type frequency, is especially scarce.

Dąbrowska (2008) tests the respective effects of type frequency and phonological similarity in an experiment that prompts speakers of Polish to add dative endings to nonce nouns. In the Polish case system, the masculine dative ending *-owi* has a higher type frequency than the feminine dative endings *-'e* and *-i*, which in turn have higher type frequencies than the neuter dative ending *-u*. Dąbrowska (2008: 937) designed stimuli in such a way that the nonce words differed systematically in their phonological neighborhood density. Items with high neighborhood density were similar to a large number of existing words, whereas items with low neighborhood density were dissimilar to other elements in the Polish lexicon. The results indicate that performance is better if the type frequency of the case inflection is higher and if the phonological neighborhood is densely populated (Dąbrowska 2008: 947).

Like high token frequency, high type frequency drives phonological reduction effects. Bybee (2002) examined final *t/d*-deletion in a corpus of spoken English, comparing pre-vocalic deletion rates across auxiliaries with contracted negation (*don't*), lexical items ending in an unstressed syllable with final *nt* (*government*, *different*), and regular past tense forms (*kissed*, *burned*). Bybee (2002: 276) finds that the highest rates of pre-vocalic *t/d*-deletion are observed with regular past tenses, which can be motivated in terms of the high type frequency of that construction. This result is corroborated by Díaz-Campos and Gradoville (2011), who study intervocalic *d*-deletion in Spanish participles with *-ado* and *-ido*. Deletion rates are higher with *-ado*, which is characterized by a relatively higher type frequency.

To summarize this section, type frequency is central for measures of constructional productivity, in which it interacts with both semantic coherence between types and phonological neighborhood density. High type frequency is furthermore correlated with rates of phonological reduction in forms that instantiate a pattern.

## 2.3 Relative frequency

Relative frequency measurements compare the token frequency of one linguistic unit to that of one or several related linguistic units and express the result in terms of a ratio. This is often done with constructions that serve as alternative expressions of similar or near-identical ideas. For example, the English dative alternation comprises two constructions, the ditransitive construction (*Mary gave John the book*) and the prepositional dative construction (*Mary gave the book to John*), which both serve to express the transfer of an object between an agent and a recipient. The constructions differ in terms of the sequence with which the arguments of the verb are presented, which means that speakers' choices between the two are commonly rooted in principles of information structure, such as the principle of end-focus (Hilpert 2021). In naturally occurring language use, the ditransitive construction has higher token frequency than the prepositional dative construction (Bresnan et al. 2007). Another pair of constructions with similar functions and a marked frequency asymmetry comprises the English active and passive. Sentences such as *The dog bit the mailman* or *The mailman was bitten by the dog* serve to express the same state of affairs with a difference in focus on either the agent or the patient. Passive sentences are much less frequent than active transitive sentences. As a third example, the English comparative can be formed morphologically and periphrastically, as in *prouder* or *more proud* (Hilpert 2008). In general English usage, morphological comparatives vastly outnumber periphrastic comparatives.

Asymmetries in relative frequency influence language use and language processing in several ways. One such effect can be observed in structural priming [see also chapter 8]. When speakers process a syntactic construction, this increases the likelihood that they will produce that construction in subsequent language use (Bock 1986). This tendency is frequency-sensitive. The so-called inverse-preference effect (Ferreira and Bock 2006) is the

phenomenon that in a pair of related constructions, structural priming is stronger for the construction with the lower relative frequency. Jaeger and Snider (2007) demonstrate this for the English dative alternation, in which prepositional datives yield a stronger priming effect, and the English voice alternation, in which passives have a more pronounced impact. The inverse-preference effect has been widely replicated. For example, Rosemeyer and Schwenter (2019) document it for the Spanish past subjunctive forms *-se* vs. *-ra* in historical corpus data. Torres Cacoullos (2015) further documents a relation between priming and the analyzability of constructions. In a study that addresses the Spanish progressive with the copula *estar* 'be' and a gerund on the basis of historical corpus data, Torres Cacoullos tests whether speakers are more likely to use the progressive construction if they have been primed with other constructions that contain the verb *estar*, such as locative, resultative, and predicate adjective constructions. The data reveal that such priming effects exist, and that they become increasingly weaker over time (Torres Cacoullos 2015: 278). This suggests that the holistic unit status of the progressive construction has become strengthened over time.

Relative frequencies are also relevant for an effect that is known as statistical preemption (Goldberg 2019). The basic premise of statistical preemption is that speakers form generalizations over pairs of constructions that serve comparable functions, such as the dative alternation or the morphological comparative and the periphrastic comparative. It is further assumed that speakers' mental representations of these construction pairs comprise knowledge of the lexical items that they encounter in each construction. For example, speakers' knowledge of the dative alternation would include the fact that verbs such as *give*, *send*, or *promise* occur both in the ditransitive construction and in the prepositional dative construction. Importantly, the relative frequency distribution of some verbs is highly asymmetric, so that for example the verb *explain* regularly occurs in the prepositional dative construction (*She explained the problem to me*), but not at all in the ditransitive construction (**She explained me the problem*). Speakers subconsciously take note of such statistical asymmetries and interpret them as grammatical constraints. In more concrete terms, this reasoning process can be spelled out as follows. Speakers know that the ditransitive construction and the prepositional dative construction convey similar meanings and occur with similar sets of verbs. They register that the ditransitive construction is overall more frequent. Yet, they encounter the verb *explain* frequently in the prepositional dative

construction, but not at all in the ditransitive construction. Since this imbalance could not be due to chance, it has to be the case that *explain* cannot be used in the ditransitive construction. In a study that approaches this issue experimentally, Boyd and Goldberg (2011) expose participants to novel adjectives such as *adax* and *ablim*. Those forms phonologically resemble English adjectives such as *awake* or *afraid*, which can be used predicatively (*The child is awake*, *The passenger is afraid*), but not attributively (**the awake child*, **the afraid passenger*). In other words, their relative frequency distribution across predicative and attributive contexts is maximally asymmetrical. Boyd and Goldberg (2011: 69) observe that the participants of their study avoided using forms like *adax* and *ablim* in attributive contexts. Statistical preemption can thus account for the fact that speakers do not produce certain expressions and find them ungrammatical, when they are asked about them. Entrenchment as such can only explain why certain expressions sound familiar to speakers, but it cannot explain the difference between previously unseen structures that are fully acceptable and unseen structures that are completely unacceptable.

Another effect of relative frequencies has been documented by Hilpert (2008) in a study of English comparative constructions. Adjectives such as *proud* or *healthy* can form the comparative either morphologically (*prouder*, *happier*) or periphrastically (*more proud*, *more healthy*). The variation between the two variants is influenced by a broad range of factors, including the length of the adjective, its stress pattern, its final phonological segment, and its syntactic context (Hilpert 2008: 407). Frequency of use impacts speakers' choices in two ways. First, there is an effect of high token frequency. A highly frequent adjective such as *easy* has a stronger tendency to form the morphological comparative than a less frequent adjective with similar phonological characteristics, such as *queasy*. A second frequency effect concerns the relative frequency of adjectives in the positive form and its comparative formations. An adjective such as *tall* is frequently used comparatively, whereas this is not the case for adjectives such as *red* or *square*, which are less gradable. Highly gradable adjectives show a greater tendency to be used with the morphological comparative.

The main point to take away from this section is the following. Effects such as statistical preemption and the inverse-preference effect in structural priming reveal that speakers are sensitive to relative frequencies in language use, and that this sensitivity impacts their choices between constructions that function as mutual alternatives.

## 2.4 Frequency of co-occurrence

Measuring frequencies of co-occurrence is one of the fundamental analytical techniques of corpus linguistics, which has a long tradition of studying collocations (Gries 2013). Collocations can be broadly defined as multi-word units that exhibit varying degrees of fixedness. They are exemplified by word pairs with non-compositional, conventionalized meanings such as *fast food*, word combinations with elements that co-occur much more frequently than would be expected by chance, such as *unmitigated disaster*, and multi-word idiomatic expressions such as *make hay while the sun shines*. The term collocation has further been used to describe the regular co-occurrence of lexical items in close proximity, which is observed for example for the words *refugee* and *crisis*.

Whereas token frequency, type frequency, and relative frequency are measured through basic counts of instances in a corpus, frequencies of co-occurrence are studied in terms of various association measures, which include for example *Mutual Information*, *Log Likelihood*, or *Delta P* (Brezina 2018: 72). These measures take into account how often two linguistic elements appear together, and how often the same elements appear on their own in other contexts. A comparison of the co-occurrence frequencies with the individual token frequencies shows how strongly the two elements of a collocation are mutually associated. A collocation such as *sustainable development* can serve as an example. Table 1, which uses data from the British National Corpus, shows the token frequency of the collocation, the token frequencies of the individual elements that make up the collocation, and the total number of words in the corpus.

| | sustainable | NOT sustainable | TOTAL |
|---|---|---|---|
| development | 147 | 31,564 | 31,711 |
| NOT development | 518 | 96,954,478 | 96,954,996 |
| TOTAL | 665 | 96,986,042 | 96,986,707 |

Table 1: A contingency table of *sustainable development* and its frequency in the BNC

The table shows that the collocation *sustainable development* appears 147 times in the corpus, which contains a total of almost 97 million words. Given the token frequencies of the individual words *sustainable* (665) and *development* (31,711), it can be computed how often the two words would be expected to appear together by pure chance. The expected frequency of the collocation is the product of the individual token frequencies, divided by the total number of words in the corpus, which in this case yields a value of 0.22. So whereas less than one instance of *sustainable development* would be expected, almost 150 are observed, which indicates a high degree of mutual association between *sustainable* and *development*. Association measures express collocation strength through scores that allow comparisons between different word pairs, so that for example the mutual attraction of *sustainable development* can be compared to that of *sustainable tourism* or *sustainable seafood*.

Whereas the study of collocations has a long tradition that focuses on associations between specific words, the family of methods known as *collostructional analysis* (Stefanowitsch and Gries 2003, Gries and Stefanowitsch 2004a, 2004b) applies the concept of collocations to the analysis of mutual associations between a construction and the lexical items that appear within it. Collostructional analysis thus takes into account the grammatical context in which lexical words appear. Stefanowitsch and Gries (2003: 219) illustrate this idea with the English *NOUN waiting to happen construction*. The construction has an initial slot for nouns that is typically filled by semantically negative words such as *disaster* or *accident*, as for examples in *It's a disaster waiting to happen* or *That was an accident waiting to happen*. In order to analyze the mutual association between the construction and the nouns that occur in it, collostructional analysis determines the token frequency of the nouns in the construction, their overall token frequency in the corpus, and the overall token frequency of the construction as such. Table 2, taken from Stefanowitsch and Gries (2003: 219), illustrates this for the noun *accident*. The lower right corner of the table represents the total of clause-level constructions in the British National Corpus. The co-occurrence frequency of *accident* and *waiting to happen* is much higher than would be expected.

| | accident | NOT accident | TOTAL |
|---|---|---|---|
| [N waiting to happen] | 14 | 21 | 35 |

| | | | |
|---|---|---|---|
| NOT [N waiting to happen] | 8,606 | 10,197,659 | 10,206,265 |
| TOTAL | 8,620 | 10,197,680 | 10,206,300 |

Table 2: A contingency table of *accident* and the [N waiting to happen] construction

The output of a collostructional analysis is a list of lexical elements that are ranked in terms of their strength of association with the construction that is being studied. Association measures that are used to obtain these rankings include the *Fisher-Yates Exact Test* and *Log Likelihood* (Flach 2021). A collostructional analysis can not only determine the lexical elements that are most strongly associated with a construction, but it also serves to identify the elements that occur significantly less often than would be expected by chance. Inspecting the words that are repelled by a construction often yields useful insights into the semantic constraints that characterize the usage of a construction.

There is no single collocation measure that could be said to reflect mutual association in the best or most objective way. Rather, measures such as *Mutual Information* or *Log Likelihood* differ in the relative importance they assign to aspects such as the token frequency of a collocation or the mutual faithfulness of the component words. For example, the collocation *unmitigated disaster* is much less frequent than *sustainable development*, but when the adjective *unmitigated* is observed, it acts as a very strong cue for the upcoming noun *disaster*. Measures such as *Delta P* further allow researchers to take the directionality of association into account. For example, the word *instead* is a strong cue that the following word would be *of*. This association is not symmetric, since *of* is frequently preceded by other elements than *instead*. Conversely, the word *vitro* is typically preceded by *in*, but *in* is not a reliable indicator that the next word will be *vitro*.

Frequencies of co-occurrence lie at the heart of distributional semantic approaches (Turney and Pantel 2010), which model the meanings of linguistic units in terms of their collocational profiles. To give a simple example, the meaning of a word such as *violin* is reflected in its distributional behavior, which is to say that words such as *piano*, *orchestra*, *play*, or *sonata* are strongly represented in contexts in which the word *violin* appears. On the basis of a corpus, it is possible to generate a frequency list of all words that appear in the neighborhood of the word *violin*. With an association measure, that frequency list can be transformed into a vector of values that shows which words are overrepresented or

underrepresented, so that a semantic profile of *violin* emerges. Distributional semantic techniques typically compare many such vectors to establish patterns of semantic similarity or dissimilarity in larger sets of linguistic units. Analyses of that kind are able to determine that, for example, words such as *piano*, *flute*, and *cello* have collocational profiles that are very similar to that of *violin*, whereas words such as *rose*, *tulip*, and *daffodil* conform to a profile that is altogether different. Perek (2018) uses a distributional semantic approach in a diachronic study of the English *way*-construction. He uses data from the Corpus of Historical American English (Davies 2010) in order to compute semantic vectors for the lexical verbs that appear in the *way*-construction. By comparing the semantic vectors of verbs that enter the construction across historical periods of time, Perek finds that the verbs of the *way*-construction occupy an increasingly broader semantic range, specifically in the manner sense of the construction (e.g. *John joked his way into the meeting*). For the path-creation sense of the construction, there is a development towards increasingly more varied and also more abstract verbs, including groups such as verbs of ingestion (*eat*, *drink*, etc.), verbs of commercial transactions (*buy*, *purchase*, *acquire*, etc.), speech act verbs (*mumble*, *whisper*, etc.), and others (*announce*, *explain*, *write*, etc.). In general terms, Perek's study illustrates how co-occurrence frequencies, put into the service of distributional semantic analyses, can reveal how constructions behave with regard to schematicity and abstractness. The elements that co-occur with a construction reflect its meaning, and historical shifts in the co-occurrence patterns between a construction and its lexical collocates can yield insights into the diachronic semantic trajectory of a construction.

Frequency of co-occurrence can be shown to affect language processing and language use. Relevant evidence is presented by Gries et al. (2005), who devised a sentence completion task in which speakers had to find continuations for sentence fragments. Gries and colleagues used the English *as*-predicative construction (*The idea was considered as a major innovation*) as a case study. The participants were exposed to sentence fragments that contained a subject, a passive auxiliary, and a past participle, as in *The idea was considered.* Their task was to complete the fragment in any way they saw fit. Gries et al. measured how often the participants' continuations resulted in a complete *as*-predicative construction. With regard to the verbs that appeared in the fragments, Gries et al. controlled for their token frequency as well as their frequency of co-occurrence with the *as*-predicative construction. For the latter, the *Fisher-Yates Exact Test* was used as an association measure

(Gries et al. 2005: 647). Verbs such as *regard*, *describe*, or *see* are strongly associated with the construction. The results of the experiment show that participants were more likely to complete a fragment with the *as*-predicative if the verb in the fragment was strongly attracted to the construction. By contrast, Gries et al. (2005: 659) did not observe an independent main effect of token frequency.

Another effect of co-occurrence frequency is observed by Hilpert and Flach (to appear) in a study of a phenomenon that is called *constructional contamination* (Pijpops and Van de Velde 2016, Pijpops et al. 2018). The term describes a relation between two constructions of a language, such that usage frequencies of one construction influence patterns of variation in another construction. Hilpert and Flach (to appear) study constructional contamination in the English passive, which exhibits variation with regard to the adverbial modification of participles. In a passive sentence, the past participle can be either preceded by an adverb, as in *The government was democratically elected*, or it can be followed by an adverb, as in *The government was elected democratically*. Pijpops and Van de Velde (2016) argue that variation of this kind can potentially be influenced by a construction that features word strings that can also appear in one of the variants. With regard to the English passive, sequences of an adverb and a participle occur in noun phrases such as *a democratically elected government*. If uses of this kind appear with high token frequency in language use, the prediction is that speakers will favor the adverb-initial variant of the passive. Hilpert and Flach (to appear) test this prediction on the basis of data from the Corpus of Contemporary American English (Davies 2008), and they find that the token frequency of adverb-participle combinations in English noun phrase constructions significantly biases speakers towards the use of adverb-initial passives. Beyond this effect of high token frequency, they further observe an effect of co-occurrence frequency. Combinations such as *dimly lit* or *randomly selected* are not highly frequent in English noun phrases. Yet, the fact that these collocations have a strong degree of mutual association may bias speakers towards adverb-initial order in the passive construction. Hilpert and Flach controlled for a possible effect of mutual association strength with covarying-collexeme analysis (Gries and Stefanowitsch 2004b), which was used to determine degrees of mutual association between adverbs and participles in the English noun phrase construction. The results indicate that frequency of co-occurrence is a significant predictor of speakers' bias towards adverb-initial order in the English passive construction.

The studies described in this section suggest that language processing and production is shaped by co-occurrence frequencies of linguistic units. The respective effects go beyond effects of high token frequency as such. Degrees of mutual association can impact language use even when the linguistic units in question are relatively low in token frequency.

## 2.5 Dispersion (burstiness)

The dispersion, or burstiness, of a linguistic unit concerns the regularity with which it appears and reappears in language use. Two linguistic units with the same token frequency may behave very different with regard to their dispersion, so that their respective tokens are either spaced out evenly and regularly, or tightly clustered together. To take an example, consider the words *without* and *system*, which appear with roughly the same token frequency in the British National Corpus. Whereas *without* appears in just about every text that is featured in the corpus, this is not the case for *system*. Tokens of the word *system* have a greater chance of following each other in close proximity, but there are also longer stretches in the corpus that do not contain the word *system* at all. What this means is that *without* is more evenly dispersed than *system*.

Various techniques have been proposed to measure dispersion (Gries 2008, 2010, 2022a). The conceptual basis for most of these techniques is that a corpus is divided into parts. The corpus parts may be defined in terms of a fixed number of running words, or by dividing the data into the different text documents that constitute the corpus. Based on such a division, it can be determined how frequently a given linguistic unit appears in each of the parts. The dispersion measure known as *range* simply assesses the percentage of corpus parts in which a linguistic unit is represented (Gries 2008: 407). The higher the range, the more even the dispersion. The dispersion measure *Deviation of Proportions* (Gries 2008: 415) is based on more sophisticated calculations. It is computed as follows. In a first step, a corpus is divided into parts, and it is determined for each part how much of the corpus it represents. A text of 10,000 words would thus represent 1% of a corpus with one million words. To assess the dispersion of a linguistic unit in that corpus, its frequency in the entire corpus and in each corpus part is determined. For example, the word *without* might appear 400 times in the corpus as a whole, and 6 times in a first text of 10,000 words. This means

that *without* is overrepresented in that corpus part. Whereas only 1% of tokens would be expected, the text in fact contains 6 tokens, which add up to 1.5%. *Deviation of proportions* is calculated in such a way that for all corpus parts, the difference between expected percentage and observed percentage are added up, and the sum of differences is divided by two. If the resulting value is close to 0, the analyzed word is very evenly distributed. Values close to 1 indicate a highly uneven dispersion.

An alternative to measuring dispersion on the basis of corpus parts is to observe distances between the tokens of a linguistic unit (Gries 2008: 409). For the word *without*, the maximal distance will be lower than for the word *London*. If the distances are distributed in such a way that some of them are very large and most are very small, this indicates an uneven dispersion.

It has been argued that certain effects of token frequency are more adequately explained as effects of dispersion. For example, Adelman et al. (2006: 817) analyzed latencies in six separate data sets from experiments that included word-naming and lexical decision tasks, finding that dispersion, operationalized as range, is a better predictor of word-processing times than token frequency, and that token frequency has no measurable effect that would be independent of dispersion. Baayen (2011: 454) draws a similar conclusion on the basis of data from lexical decision tasks, arguing that any effect of token frequency is minimal once variables such as dispersion, morphological family size, or syntactic entropy are controlled for. It is important to note in this context that most of the established measures of dispersion, including range, are known to correlate strongly with token frequency (Adelman et al. 2006: 815, Baayen 2011: 445). Infrequent lexical items will necessarily appear in only few corpus parts, whereas frequent items stand a better chance of being represented in a large percentage of the parts. In order to address this problem, Gries (2022a) has recently proposed a measure that avoids this correlation and thus allows for a more reliable assessment of the respective effects of token frequency and dispersion. Gries uses *Deviation of proportions* as the basis for a new measure that is labeled $DP_{nofreq}$. The measure calculates theoretical lower and upper bounds for the dispersion of a linguistic item, given its token frequency. In other words, the measure assesses whether a word could be potentially more even or more uneven in its dispersion. Unlike other dispersion measures, $DP_{nofreq}$ can yield high dispersion values for infrequent linguistic items. Gries (2022a: 29) uses data from lexical decision tasks to show that $DP_{nofreq}$, in combination with

the variables of token frequency and word length, outperforms other dispersion measures in predicting reaction times.

Besides facilitating greater ease of processing, evenness of dispersion also correlates with aspects of linguistic meaning. Using a distance-based measure of dispersion, Pierrehumbert (2012: 104) shows that linguistic units with more specialized meanings are dispersed less evenly. Specifically, in a comparison of deverbal nouns (*evolution*, *failure*, *growth*) and their verbal bases (*evolve*, *fail*, *grow*), the nouns show a stronger tendency to appear in bursts, whereas the verbs are dispersed more evenly. In a study with a similar outlook, Hilpert and Correia Saavedra (2017) analyze a matched set of lexical words and grammatical items that are balanced in terms of their respective token frequencies. They use regression modeling to determine whether even dispersion, measured in terms of *deviation of proportions* (Gries 2008), is predicted more accurately by high token frequency or by semantic generality. The results indicate that both token frequency and semantic generality have significant effects. The semantic effect is however considerably weaker than the frequency effect.

In comparison to other measures of frequency, dispersion has received relatively less attention in research that is concerned with constructions and Cognitive Linguistics more generally. The importance of considering dispersion and its effects has been argued forcefully by Gries in a series of papers (2008, 2010, 2022a, 2022b). Aside from the issues that have been presented in this section, Gries (2022b: 62) further suggests that the effects of dispersion on language learning merit further attention.

## 3 Conclusions and outlook

The discussion in this chapter up to this point allows three broad conclusions. First, the notion of frequency has become increasingly central to research in usage-based Construction Grammar. This development is not only reflected in the definitions that Goldberg (1995, 2006, 2019) offers for the term construction, but also in a steadily growing number of studies that present their arguments on the basis of corpus-based frequency data. Second, frequency is considered in that work not just as token frequency, but as a range of several different measures. This chapter has reviewed studies drawing on token

frequency, type frequency, relative frequency, co-occurrence frequency, and dispersion. Section 2 covered the ways in which these aspects of frequency are measured, and to what ends they are being analyzed. The third conclusion is that frequency is intimately related to issues of language processing and production. The discussion has touched on a wide variety of frequency effects, including the relation of high token frequency and entrenchment (Ellis 2002: 152), the inverse-preference effect that links relative frequency and structural priming (Ferreira and Bock 2006), and the relation of uneven dispersion and specialized meaning (Pierrehumbert 2012: 104). Understanding how these frequency effects work is important for theoretical work in Construction Grammar, notably with regard to the organization of constructions in a network (Diessel 2019, Schmid 2020).

It is without dispute that the increasingly thorough engagement with frequency-related issues has yielded useful insights and has thereby advanced Construction Grammar as a field. That said, there are also reasons to maintain a critical view of what has been accomplished and what still remains to be done. The remainder of this section will go over four points that merit attention. First, Blumenthal-Dramé (2017: 141) points out that "*we are still a long way from fully understanding the intricate relationships among usage frequency, entrenchment, and others factors that might modulate the strength and autonomy of linguistic representations in our minds*". How exactly the experience of a string of words impacts the mental representation of a schematic construction that is instantiated by those words remains to be worked out, especially with regard to the interplay of token frequency, type frequency and co-occurrence frequency. Second, it is crucial not to attribute effects to frequency without considering alternative explanations. Section 2.5 pointed to work that tested systematically whether the effects of token frequency persist if other frequency-related factors such as dispersion are included. Gries (2022b: 47) cautions against the reliance on token frequency measurements that do not take variability across corpora and corpus parts into account. Variability across and within corpora can be substantial, but work that takes this variability into account remains the exception. Third, the aspects of frequency that were discussed in this paper have not received the same kind of attention. Whereas measurements of token frequency and co-occurrence frequency are routinely applied in constructional analyses, measures of dispersion are rarely taken into account (Gries 2008: 403), and in most studies that consider it, the measures that are used are strongly confounded with token frequency (Gries 2022a). What this means is that more work

is necessary in order to realistically assess the effects of dispersion and the interplay between dispersion and other aspects of frequency. The fourth and final point is of a more general nature. Whereas this chapter has laid out a range of different aspects of frequency, measures that are intended to capture these aspects, and effects on processing and production that stem from them, it is important to recognize that any linguistic phenomenon is likely to benefit from an analysis that takes these different aspects into account simultaneously, in order to capture how different frequency effects interact. For example, in order to arrive at a comprehensive understanding of an argument structure construction such as the caused motion construction (Goldberg 1995), several different frequency measures would be useful, including token frequencies of the construction as such, token and type frequencies of the elements that appear in its slots, co-occurrence frequencies of the construction and lexical elements, relative frequencies of constructions that are similar in form or function, and dispersion measures that assess the distribution of the construction across different parts of the used corpora. It is clear that comprehensive analyses of this kind are laborious, but they hold the promise of coming to terms with the multi-faceted role of frequency in language use, which is an important goal for the future development of Construction Grammar.

**References**


Adelman, James, Gordon D.A. Brown, and José F. Quesada. 2006. Contextual diversity, not word frequency, determines word-naming and lexical decision times. *Psychological Science*, 17/9, 814.

Arnon, Inbal and Neal Snider. 2010. More than words: Frequency effects for multi-word phrases. *Journal of Memory and Language* 62/1: 67–82.

Baayen, R. Harald. 2009. Corpus linguistics in morphology: morphological productivity. In: Anke Lüdeling and Merja Kytö (eds.), *Corpus Linguistics. An international handbook*. Berlin: Mouton De Gruyter, 900-919.

Baayen, R. Harald. 2011. Demythologizing the word frequency effect: A discriminative learning perspective. *The Mental Lexicon* 5, 436–461.

Barðdal, Jóhanna. 2008. *Productivity: Evidence from Case and Argument Structure in Icelandic*. Amsterdam: John Benjamins.

Blumenthal-Dramé, Alice. 2017. Entrenchment from a Psycholinguistic and Neurolinguistic Perspective. In Hans-Jörg Schmid (ed.), *Entrenchment and the Psychology of Language Learning: How We Reorganize and Adapt Linguistic Knowledge*, Boston: APA and Walter de Gruyter, 129–52.

Bock, J. Kathryn. 1986. Syntactic persistence in language production. *Cognitive Psychology* 18: 355–87.

Boyd, Jeremy K. and Adele E. Goldberg. 2011. Learning what not to say: The role of statistical preemption and categorization in "a"-adjective production. *Language* 81/1: 1–29.

Brezina, Vaclav. 2018. *Statistics in Corpus Linguistics. A practical guide*. Cambridge: Cambridge University Press.

Bresnan, Joan, Anna Cueni, Tatiana Nikitina, and Harald R. Baayen. 2007. Predicting the dative alternation. In G. Boume, I. Kraemer, and J. Zwarts (eds.) *Cognitive Foundations of Interpretation*. Amsterdam: Royal Netherlands Academy of Science, 69–94.

Brooks, Patricia and Michael Tomasello. 1999. Young children learn to produce passives with nonce verbs. *Developmental Psychology* 35: 29–44.

Brooks, Patricia, Michael Tomasello, Kelly Dodson and Lawrence B. Lewis. 1999. Young children's overgeneralizations with fixed transitivity verbs. *Child Development* 70, 1325–1337.

Bybee, Joan. 1995. Regular morphology and the lexicon. *Language And Cognitive Processes* 10, 425–455.

Bybee, Joan L. 2002. Word frequency and context of use in the lexical diffusion of phonetically conditioned sound change. *Language Variation and Change* 14, 261-90.

Bybee, Joan. 2006. From usage to grammar: the mind's response to repetition. *Language* 82/4, 711-733.

Bybee, Joan. 2010. *Language, Usage and Cognition*. Cambridge: Cambridge University Press.

Bybee, Joan and Joanne Scheibman. 1999. The effect of usage on degrees of constituency: the reduction of *don't* in English. *Linguistics* 37/4, 575-596.

Bybee, Joan and Sandra Thompson. 1997. Three frequency effects in syntax. *Berkeley Linguistic Society* 23, 65-85.

Clark, Eve V. 1987. The principle of contrast: A constraint on language acquisition. In B. MacWhinney (ed.), *Mechanisms of language acquisition*. Hillsdale: Lawrence Erlbaum, 1-33.

Dąbrowska, Ewa. 2008. The effects of frequency and neighbourhood density on adult speakers - productivity with Polish case inflections: An empirical test of usage-based approaches to morphology. *Journal of Memory and Language* 58/4, 931–951.

Davies, Mark. 2008. The Corpus of Contemporary American English (COCA): 400+ million words, 1990-present. Available online at http://www.americancorpus.org.

Davies, Mark. 2010. The Corpus of Historical American English (COHA): 400+ Million Words, 1810–2009, available online at http://corpus.byu.edu/coha.

Díaz-Campos, Manuel and Michael Gradoville. 2011. An analysis of frequency as a factor contributing to the diffusion of variable phenomena: Evidence from Spanish data. In . Luis A. Ortiz-López (ed.), *Selected Proceedings of the 13th Hispanic Linguistics Symposium*. Somerville, MA: Cascadilla Proceedings Project, 224-238.

Diessel, Holger. 2019. *The Grammar Network: How Linguistic Structure Is Shaped by Language Use*. Cambridge: Cambridge University Press.

Ellis, Nick C. 2002. Frequency effects in language processing: A review with implications for theories of implicit and explicit language acquisition. *Studies in Second Language Acquisition* 24, 143-188.

Ferreira, Victor S. and Kathryn J. Bock. 2006. The Functions of Structural Priming. *Language and Cognitive Processes* 21, 1011-1029.

Flach, Susanne. 2021. Collostructions: An R implementation for the family of collostructional methods. Package version v.0.2.0, https://sfla.ch/collostructions/.

Fraser, Bruce. 2009. Topic Orientation Markers. *Journal of Pragmatics* 41/5, 892-898.

Goldberg, Adele E. 1995. *Constructions: A Construction Grammar Approach to Argument Structure*, Chicago: University of Chicago Press.

Goldberg, Adele E. 2006. *Constructions at Work: The Nature of Generalization in Language*, Oxford: Oxford University Press.

Goldberg, Adele E. 2019. *Explain Me This: Creativity, Competition, and the Partial Productivity of Constructions*. Princeton: Princeton University Press.

Gries, Stefan Th. 2008. Dispersions and adjusted frequencies in corpora. *International Journal of Corpus Linguistics* 13/4, 403-437.

Gries, Stefan Th. 2010. Dispersions and adjusted frequencies in corpora: further explorations. In Stefan Th. Gries, Stefanie Wulff, and Mark Davies (eds.), *Corpus linguistic applications: current studies, new directions*. Amsterdam: Rodopi, 197-212.

Gries, Stefan Th. 2013. 50-something years of work on collocations: what is or should be next … *International Journal of Corpus Linguistics* 18/1. 137-165.

Gries, Stefan Th. 2022a. What do (most of) our dispersion measures measure (most)? Dispersion? *Journal of Second Language Studies*.

Gries, Stefan Th. 2022b. On, or against?, (just) frequency. In Hans C. Boas (ed.), *Directions for pedagogical Construction Grammar: Learning and teaching (with) constructions*. Berlin: De Gruyter, 47-71.

Gries, Stefan Th. and Dagmar S. Divjak (eds.). 2012. *Frequency effects in language learning and processing*. Berlin: De Gruyter Mouton.

Gries, Stefan Th., Beate Hampe, and Doris Schönefeld. 2005. Converging evidence: Bringing together experimental and corpus data on the association of verbs and constructions, *Cognitive Linguistics* 16-4, 635-676.

Gries, Stefan T. and Anatol Stefanowitsch. 2004a. Extending collostructional analysis: A corpus-based perspective on “alternations”. *International Journal of Corpus Linguistics* 9/1: 97–129.

Gries, Stefan T. and Anatol Stefanowitsch. 2004b. Co-varying collexemes in the into-causative. In Michel Achard and Suzanne Kemmer (eds), *Language, Culture, and Mind*, Stanford: CSLI, 225–36.

Hay, Jennifer and Paul Foulkes. 2016. The evolution of medial /t/ over real and remembered time. *Language* 92/2, 298-330.

Hilpert, Martin. 2008. The English comparative - language structure and language use. *English Language and Linguistics* 12/3, 395-417.

Hilpert, Martin. 2015. Historical Linguistics. In Ewa Dabrowska and Dagmar Divjak (eds.), *Handbook of Cognitive Linguistics*. Berlin: De Gruyter Mouton, 346-365.

Hilpert, Martin. 2021. Information structure. In Bas Aarts, April MacMahon, and Lars Hinrichs (eds.), *Handbook of English Linguistics*. New York: Wiley, 229-248.

Hilpert, Martin and David Correia Saavedra. 2017a. Why are grammatical elements more evenly dispersed than lexical elements? Assessing the roles of frequency and semantic generality. *Corpora* 12(3). 369-392.

Hilpert, Martin and Susanne Flach. to appear. A case of constructional contamination in English: Modified noun phrases influence adverb placement in the passive.

Hoffmann, Thomas. 2019. *English comparative correlatives: Diachronic and synchronic variation at the lexicon-syntax interface*. Cambridge: Cambridge University Press.

Hopper, Paul J. and Elizabeth C. Traugott. 2003. *Grammaticalization*. 2nd ed. Cambridge: Cambridge University Press.

Israel, Michael. 1996. The way constructions grow. In Adele Goldberg (ed.), *Conceptual Structure, Discourse and Language*. Stanford: CSLI, 217–30.

Jaeger, T. Florian and Neal Snider. 2007. Implicit learning and syntactic persistence: Surprisal and cumulativity. In L. Wolter and J. Thorson (eds.), *University of Rochester Working Papers in the Language Sciences* 3(1), 26-44.

Jucker, Andreas H. and Sara W. Smith. 1998. And people just you know like ‘wow’: Discourse markers as negotiating strategies. In: Jucker, Andreas and Yael Ziv (eds.), *Discourse Markers: Descriptions and Theory*. John Benjamins, Philadelphia, 171–201.

Kay, Paul and Charles J. Fillmore. 1999. Grammatical constructions and linguistic generalizations: The What's X Doing Y? construction. *Language* 75/1: 1–33.

Lorenz, David. 2013. *Semi-Modal Contractions in English: The Emancipating Effect of Frequency*. Freiburg: Rombach.

Perek, Florent. 2018. Recent change in the productivity and schematicity of the way-construction: a distributional semantic analysis. *Corpus Linguistic and Linguistic Theory*, 14/1, 65-97.

Pfänder, Stefan and Behrens, Heike. 2016. Experience counts. An introduction to frequency effects in language. In Behrens and Pfänder (eds.), *Experience Counts: Frequency Effects in Language*. Berlin: de Gruyter, 1-20.

Pierrehumbert, Janet B. 2012. Burstiness of Verbs and Derived Nouns. In Diana Santos, Krister Linden, and Wanjiju Ng'ang'a (eds.), *Shall we Play the Festschrift Game?: Essays on the Occasion of Lauri Carlson's 60th Birthday*. Berlin: Springer, 99-116.

Pijpops, Dirk and Freek Van de Velde. (2016). Constructional contamination: How does it work and how do we measure it? *Folia Linguistica* 50/2, 543-582.

Pijpops, Dirk, Isabeau de Smet, and Freek Van de Velde. 2018. Constructional contamination in morphology and syntax: four case studies. *Constructions and Frames* 10/2, 269-305.

Rivas, Javier and Esther L. Brown. 2012. Stage-Level and Individual-Level distinction in morphological variation. *Borealis – An International Journal of Hispanic Linguistics* 1/2, 73–90.

Rosemeyer, Malte and Scott E. Schwenter. 2019. Entrenchment and persistence in language change: the Spanish past subjunctive. *Corpus Linguistics and Linguistic Theory* 15/1, 167-204.

Schmid, Hans-Jörg. 2020. *The dynamics of the linguistic system. Usage, conventionalization, and entrenchment*. Oxford: Oxford University Press.

Stefanowitsch, Anatol and Stefan Th. Gries. 2003. Collostructions: Investigating the interaction between words and constructions. *International Journal of Corpus Linguistics* 8/2, 209-43.

Suttle, Laura and Adele E. Goldberg. 2011. The partial productivity of constructions as induction. *Linguistics* 49/6, 1237-1269.

Torres Cacoullos, Rena. 2015. Gradual loss of analyzability: Diachronic priming effects. In Aria Adli, Marco García García and Göz Kaufmann (eds.), *Variation in Language: System- and Usage-based Approaches*. Berlin: De Gruyter, 265-288.

Tottie, Gunnel. 1991. Lexical diffusion in syntactic change: Frequency as a determinant of linguistic conservatism in the development of negation in English. In Dieter Kastovsky (ed.), *Historical English Syntax*. Berlin, Germany: Mouton de Gruyter, 439-467.

Turney, Peter D. and Patrick Pantel. 2010. From Frequency to Meaning: Vector Space Models of Semantics. *Journal of Artificial Intelligence Research* 37, 141-188.

Yoon, Jiyoung and Stefan Th. Gries (eds.). 2016. *Corpus-based Approaches to Construction Grammar*. Amsterdam: John Benjamins.