

Introduction

1.1 Overview: Argument realization

All languages provide ways to talk about events and their participants; this function is typically assumed in great part by verbs. It is precisely for this reason that, more so than other content words, verbs are rarely uttered in isolation but are usually accompanied by certain other words, called the *arguments* of the verb. This book is concerned with the topic of argument realization, i.e., that part of the grammar that determines how participants to verbal events are expressed in the clause.¹

For example, the verb *kill* is typically used with reference to at least two arguments, a killer and a victim. Speakers of English must know that each of these arguments is identified by a specific position in the clause: in the canonical word order, the killer is realized in the pre-verbal position (called the subject in traditional grammar), while the victim is realized in the post-verbal position (the direct object), as in *Brutus killed Caesar*. This knowledge constitutes the *argument structure* of the verb *kill*. By contrast, there are other two-place predicates that behave differently. For instance, the verb *stare* is commonly used in conjunction with two arguments, an observer and a target. The latter of these arguments is not realized as a direct object but as a prepositional phrase headed by *at*, as shown by (1a) vs. (1b) below.

- (1) a. *He stared me.
 b. He stared at me.

The observations made so far may give the impression that argument structure trivially consists of knowledge tied to individual verbs. However, the following examples from the Corpus of Contemporary American English (COCA; Davies, 2008) do not fit well with this view.

1. It should be duly noted at the outset that the terms “argument realization” and “argument structure” are not restricted to verbs, but can *a priori* apply to any kind of word that can be seen as “governing” other elements in the sentence. It is, however, verbs that have received the most attention in the literature on argument realization, probably because they can be seen, as Levin & Rappaport Hovav (2005, p. 33) put it, as “the prototypical predicators, that is, argument-taking words”. Some studies also apply the related term of “valency” to other parts of speech, such as nouns, adjectives, and prepositions (Herbst, 1983; Herbst & Schüller, 2008).

- (2) a. He stared her into immobility.
- b. Chess coughed smoke out of his lungs.
- c. [H]er nose was so bloodied that the ref whistled her off the floor.
- d. Navin sneezed blue pollen onto his shirt.

These examples strike us as rather creative. What they have in common is that they feature an unusual use of each of the verbs. As previously mentioned, *stare* usually takes two arguments and describes an event of looking, but its use in (2a) with a third argument (*into immobility*) also conveys the causation of a change of state. Similarly, the verbs *cough*, *whistle*, and *sneeze* typically refer to actions involving a single individual, and we certainly do not want to claim that they inherently convey the idea that these actions may result in the motion of some external entity (*smoke/her/blue pollen*), as in (2b–d).

In the face of such examples, the idea that argument structure is primarily knowledge about verbs loses some of its appeal. Speakers are able to use verbs in new, creative ways to fit their communicative needs, which points to broader principles of argument realization. As a matter of fact, such principles are necessary to explain certain facts about children's early use of language, in particular their *overgeneralization errors* (Baker, 1979), as exemplified by (3a) and (3b) below (from Bowerman, 1982a).

- (3) a. I'm just gonna fall this on her.
- b. I disappeared a bear in the back of the car.

These two naturally occurring sentences uttered by preschoolers involve the combination of a verb (*fall* and *disappear*) with an argument structure that is not acceptable in adult usage. Since the children cannot have heard these odd combinations from their caregivers, they must have applied a generalization beyond what is normally acceptable. They would not be able to do so if they had only learned facts about individual verbs.

It is sometimes difficult to decide on purely intuitive grounds whether a given use of an argument structure can be attributed to the main verb alone, or whether it is better captured as a creative application of a generalization. Such problems will be the central concern of Part I of this book. For example, many transitive verbs in English can be used with an indirect object to refer to the intended recipient of the result of the action described by the verb, as exemplified by (4) and (5) below (also from COCA).

- (4) a. Ask the butcher to grind you a little pork.
- b. Jerry lit us a candle from the emergency kit.
- c. I'll just chop you a bit of wood.
- (5) a. His mother cooked us an Indian-style meal.
- b. Can I buy him a cookie?
- c. I'll write you a letter every day.

It is unproblematic to claim that the sentences under (4) are creative uses, since the verbs *grind*, *light*, and *chop* do not in themselves make reference to an intended recipient or beneficiary (*you* in [4a] and [4c], *us* in [4b]), and there is nothing in the meaning of these verbs that could predispose them to occur with one. The sentences under (5) could be analyzed along similar lines, although the term “creative” seems less apt. Certainly, *cook*, *buy*, and *write* do not necessarily involve an intended recipient of, respectively, the prepared food, the purchased goods, or the written material. However, preparing meals for other people, purchasing items for someone, or writing some text (a letter, a prescription, a recommendation, etc.) intended for somebody else, are common occurrences, and the corresponding formulations “cook/buy/write someone something” will most likely be familiar to speakers of English. Tallying such uses with those under (4) fails to capture this perceived conventionality. A similar problem arises more generally with optional arguments. For example, the verbs *rent* and *sell* arguably presuppose a recipient, which can be specified as an indirect object, but can also be omitted, in which case the recipient argument receives a generic or indefinite interpretation:

- (6) a. She rented (them) the apartment.
 b. They sold (us) the house.

If there is some mechanism that allows speakers of English to use transitive verbs with an additional recipient argument, and since the transitive use of *rent* and *sell* is possible and does not produce a sense of “incompleteness”, why would it not be assumed that *rent* and *sell* are essentially transitive and that their occurrence with a recipient argument is licensed by the same mechanism that produces (4) and (5)? Linguists have traditionally relied on their intuitions to decide such matters. I will suggest an alternative approach, based on linguistic usage (cf. Chapters 2 and 3).

The examples discussed so far also illustrate another important property of argument structure, i.e., that it produces meaning. Verbs can typically occur with more than one argument structure, with systematic variations in meaning. For instance, using a transitive verb with an indirect object, as in (4) and (5), introduces the idea of an intended transfer. The following examples from Goldberg (1995, p. 11) illustrate how different argument structures introduce semantic differences between uses of the same verb, *kick*:

- (7) a. The horse kicks.
 b. Pat kicked the ball.
 c. Pat kicked at the football.
 d. Pat kicked Bob the ball.
 e. Pat kicked the football into the stadium.
 f. Pat kicked Bob black and blue.

These sentences share the notion that the agent performs some gesture, namely a forward motion of the foot, but they differ in many other respects. In (7a), the horse acts alone, in (7b) and (7c), Pat's action is targeted at a specific object, and in (7d), it causes Bob to receive the ball. In (7e), Pat's action on the football causes it to move along a defined trajectory, and in (7f), Pat's action on Bob imparts a change of state on the latter.

Trivially, sentences with a different number of arguments are likely to differ in meaning, to the extent that the semantic interpretation must be adapted to accommodate additional arguments. Yet, even pairs of sentences with the same number and type of arguments also exhibit meaning differences, as seen in the contrast between (7b) and (7c). In the most likely scenario, (7b) refers to a situation in which John kicks and makes contact with the ball, which causes it to move. However, (7c) does not necessarily imply that the ball was set in motion, neither that contact was made; rather, it merely describes an attempt by John to cause the ball to move (cf. Levin, 1993, pp. 64–65).

While descriptive accounts of the semantic aspects of argument structure abound (Jackendoff, 1983; Dixon, 1991; Rappaport Hovav & Levin, 1998; Pinker, 1989; Goldberg, 1995, *inter alia*), linguists have more rarely addressed the question of how argument structure comes to be associated with meaning in the first place, and what determines this meaning. Part II will be concerned with this question. One popular view is that argument structures derive their meaning from the verbs that frequently occur in them (cf. Goldberg, Casenhiser, & Sethuraman, 2004; Stefanowitsch & Gries, 2003). For instance, the notion of transfer conveyed by the ditransitive argument structure is taken to originate in its frequent occurrence with *give* and other verbs of giving. While this view seems adequate for argument structures that have a relatively concrete and stable meaning, it is less apt at capturing cases where the semantic contribution is more abstract and variable (cf. Chapter 4). As already mentioned, the insertion of *at* after the transitive verb *kick* usually produces an interpretation where contact with the direct object referent is not made. For one thing, it is not clear what verb(s) the meaning 'lack of contact' could stem from. Also, there are similar pairs of sentences with other verbs that do not display the same semantic contrast, as exemplified with *pull* in (8a) vs. (8b) below. Both sentences certainly entail that Bill made contact with the lever; the difference between them lies in that (8a) but not (8b) entails that the lever was successfully moved.

- (8) a. Bill pulled the lever.
 b. Bill pulled at the lever.

This issue relates to the question of what level of generalization best captures argument structure: how can the structures used in (7c) and (8b) be considered the same if they do not constitute a consistent semantic generalization? Can a single overarching

construction be maintained, or is it preferable to posit several low-level generalizations? In Chapter 5, I present usage-based evidence for the latter of these alternatives.

Finally, a last aspect that has received some attention in the literature is how different argument structures may be related. For instance, many verbs conventionally used in the double-object pattern exemplified by (4) and (5) may also occur with the same number and type of arguments in another construction with *to*, as exemplified by (9a) and (9b) below.

- (9) a. He gave/lent/promised/offered/sent/sold Larry a book.
b. He gave/lent/promised/offered/sent/sold a book to Larry.

The fact that many of such sentence pairs involve little variation in meaning (if any) might lead us to consider the two patterns as grammatical variants, and suggest that one could be derived from the other (and possibly vice versa), or at least related to it in some way. On the other hand, each variant is associated with particular constraints that restricts its use (cf. Erteschik-Shir, 1979; Thompson, 1990; Bresnan, Cueni, Nikitina, & Baayen, 2007). For instance, the double-object variant is unacceptable with lexical recipients and pronominal themes (at least in American English), ruling out sentences such as **He gave Larry it*. Hence, the two variants are not always interchangeable, which suggests that they should be regarded as separate patterns. Similar observations can be shown to hold to varying extents for many other pairs of argument structures (see Levin, 1993 for a number of potential examples), which could be analyzed along the same lines.

Theories of argument structure have favored either one or the other position, although there is arguably some merit in both: the former position explicitly captures the relatedness between variants (e.g., Jackendoff, 1975; Pinker, 1989), while the latter emphasizes the specific function of each (e.g., Goldberg, 1995, 2002; Michaelis & Ruppenhofer, 2001). This begs the question of whether these two seemingly opposite views could in fact be reconciled (cf. Chapter 6). A more interesting question, however, is whether both are needed to account for the linguistic behavior of speakers (cf. Chapter 7). For instance, to what extent does the occurrence of a verb in one of two related variants trigger the expectation that it can also be used in the other variant, and how adequately do the two perspectives capture that behavior?

This book examines argument realization from the perspective of a usage-based approach to grammar, i.e., under the assumption that crucial aspects of grammatical organization are tied to the frequency with which particular words and syntactic structures are used. More particularly, it addresses the issues sketched above by appealing to usage-based explanations. In the next section, I briefly introduce the theoretical framework of usage-based linguistics and its basic tenets, and I motivate its use for the study of argument realization.

1.2 Usage-based linguistics

Usage-based theory takes the view that the cognitive representation of language emerges through, and is shaped by, language use (Langacker, 1987, 2000; Hopper, 1987; Barlow & Kemmer, 2000; Bybee, 2006, 2010, 2013; Bybee & Hopper, 2001; Bybee & McClelland, 2005). Hence, usage-based approaches reject the strict separation of grammar (or *competence*) and usage (or *performance*) typical of generative approaches, and rather than construing grammar as a mere repository to be accessed in language use, they take the view that grammar is itself the product of usage. In Bybee's (2006, p. 1) words, "grammar [is] the cognitive organization of one's experience with language". Methodologically, it follows from this approach that one can gain insights into the language system by analyzing usage data.

In usage-based approaches, grammar is commonly seen as a vast inventory of symbolic conventions that are extracted from full-fledged utterances through a gradual process of schematization, retaining the syntactic and semantic commonalities across different usage events. In line with this account of the emergence of grammar, another important point of departure from generative approaches is the exposure of the so-called "rule/list fallacy" (Langacker, 1987, p. 29), whereby linguistic units must either be produced by maximally abstract rules of grammar, or be listed in a lexicon containing all the irregularities. Instead, in a usage-based approach, grammatical patterns may be defined at any level of abstraction. Specific instances and the structures abstracted from them can be stored simultaneously; in other words, lower-level units need not be discarded once a higher-level generalization over these units has been made. In line with current thinking on categorization in cognitive psychology (Medin & Schaffer, 1978; Nosofsky, 1986, 1988; Nosofsky, Pothos, & Wills, 2011, *inter alia*), some recent versions of the usage-based approach (notably Pierrehumbert's [2001] and Bybee's [2006, 2010, 2013]) propose that the cognitive representation of language essentially consists of a large number of exemplars stored in all their details in the course of language use. In these models, abstract structures tend to be viewed as an emergent property of exemplar storage rather than as explicit rules and schemas.

While the core tenets of usage-based linguistics can in theory be understood independently of particular grammatical frameworks, they are in practice susceptible to displaying varying degrees of compatibility with specific models. Functionalist approaches (e.g., Givón, 1984, 1990; Dik, 1989; Halliday, 1994; Van Valin & LaPolla, 1997) will generally tend to have more affinity with a usage-based conception of grammar than generative frameworks (e.g., Chomsky, 1965, 1981, 1995; Borer, 2003; Ramchand, 2008; Randall, 2010). Due to their theoretical kinship with usage-based linguistics and some shared assumptions about the nature of language and its relation to the rest of cognition, cognitive-linguistic approaches

to grammar and in particular the various strands of *construction grammar* (Fried & Östman, 2004; Kay & Fillmore, 1999; Lakoff, 1987) are highly compatible with usage-based theory. In fact, several constructionist frameworks explicitly adopt a usage-based approach, such as Langacker's (1987, 1991, 2000, 2008) Cognitive Grammar, Goldberg's (1995, 2006) Cognitive Construction Grammar, and Croft's (2001) Radical Construction Grammar. In the present study, I will also adopt a constructional approach. As I argue in Chapter 2, the concept of construction allows a better account of the facts of argument realization (especially regarding its verb-general aspects), which provides further motivation for adopting an approach along the lines of construction grammar.

A wealth of evidence has been accumulated over the past few decades in support of a usage-based view of language. Many findings indicate that frequency of occurrence appears to be an important factor in linguistic representations (cf. the reviews by Ellis [2002] and Diessel [2007]). Frequent words tend to be phonologically reduced; for example, Bybee (2000) finds that deletion of final /t,d/ in American English is more common in highly frequent forms (e.g., *can't*, *don't*) than in less frequent ones, and Gahl (2008) reports that vowel duration tends to be shorter in the more frequent member of homophone pairs (e.g., *time* and *thyme*) than in the less frequent one. Frequent words are also more resistant to morphosyntactic change (Bybee, 1985, 1995; Bybee & Slobin, 1982; Hooper, 1976). Thus, high frequency is one of the factors that may lead to the formation of morphologically irregular forms, like for instance the irregular English past tense forms (e.g., *brought*, *took*, *went*, etc.). This is congruent with evidence that frequent inflected forms are stored and retrieved as whole units and not computed "on the fly", even if they are fully regular (Stemberger & MacWhinney, 1986, 1988; Hare, Ford, & Marslen-Wilson, 2001); hence, such forms may endure and outlive the word formation process that licensed them if the latter loses productivity and falls into disuse. Along similar lines, behavioral and neurological evidence indicates that the tendency of morphologically complex words (e.g., *happiness*, *insane*) to be stored and retrieved as a whole (as opposed to being compositionally derived, for instance by affixation) correlates with the frequency of the complex form (e.g., *childish*) relative to the frequency of the base form (e.g., *child*) (Hay, 2001; Blumenthal-Dramé, 2012).

Similar frequency effects are also reported at the phrasal level. Frequent sequences of words (e.g., *all over the place*, *don't have to worry*, *you don't want to*) are processed more easily (Tremblay, Derwing, & Libben, 2009; Arnon & Snider, 2010) and are repeated faster and with fewer mistakes (Bannard & Matthews, 2008; Bod, 1998; Tremblay, Derwing, Libben, & Westbury, 2011) than infrequent ones. Bybee & Scheibman (1999) also report that the phonological reduction of *don't* mostly occurs in high-frequency phrases such as *I don't know*. By the same token, Jurafsky, Bell, Gregory, & Raymond (2001) show that phonological reduction is more likely

when the conditional probability between two successive words is particularly high (see also Gregory, Raymond, Bell, Fosler-Lussier, & Jurafsky, 1999). Hilpert (2008) finds a corresponding effect of co-occurrence frequency on speech perception: when hearing words rendered phonologically ambiguous by manipulation (e.g., *cry* altered to sound more like *try*), speakers perceive the form that is the most likely given the syntactic context. At the clausal level, there is a wealth of psycholinguistic evidence that sentences tend to be processed more easily when individual verbs are used with complements that are statistically more likely (Garnsey, Pearlmutter, Myers, & Lotocky, 1997; Wiechmann, 2008; Wilson & Garnsey, 2009). Similar effects are reported in other domains of syntax, such as noun phrase conjunction (Desmet & Gibson, 2003) and relative clause attachment (Desmet, De Baecke, Drieghe, Brysbaert, & Vonk, 2006; Reali & Christiansen, 2007). Gahl & Garnsey (2004) also find that verbs taking either a direct object (e.g., *I heard the story yesterday*) or a sentence complement (e.g., *I heard the story was true*) are pronounced faster and with more final /t,d/ deletions when the post-verbal constituent matches the probabilistic syntactic preference of the verb as measured by usage frequency. Tily et al. (2009) report a similar finding with dative verbs alternating between double-object (e.g., *that gives them the full right to test the kids*) and prepositional constructions (e.g., *I haven't given much thought to it*), for which they find that the post-verbal complements are produced faster and with fewer disfluencies (filled pauses, hesitations, repairs, etc.) when they match the construction that is more likely according to a range of frequency-based factors (see also Kuperman & Bresnan, 2012). Finally, frequent syntactic patterns also tend to be more resistant to change (Bybee & Thompson, 1997). For example, Ogura (1993) observes that the spread of *do*-support in questions (e.g., *do you like...?* instead of *like you...?*) and negative sentences (e.g., *you do not like... instead of you like not...*) was resisted the longest by high-frequency verbs such as *say*, *think*, and *know*. By the same token, Bybee (2010) argues that the syntactic properties of auxiliary verbs (*be*, *have*, and the modal auxiliaries *can*, *must*, *may*, *will*, etc.) in Present-Day English are largely explained by the high frequency of occurrence of these verbs in the older question and negation constructions without *do*-support (viz. “*can/must/will... SUBJECT VERB ...?*” and “*SUBJECT can/must/will... not VERB*”).

All of these findings, and many more, are not easily captured by an approach that maintains a unidirectional relation from grammar to usage. By contrast, such effects are naturally predicted by a usage-based account, and actually provide evidence for one important corollary of the usage-based hypothesis, namely that differences in frequency should result in differences in linguistic structure. The impact of frequency on linguistic representations is often discussed in terms of *entrenchment*, which Langacker (2000, p. 3) defines as the degree to which “a highly complex event can coalesce into a well-rehearsed routine that is easily elicited and reliably executed”.

The degree of entrenchment of a linguistic structure is related to its frequency of use, and is commonly seen to bear on its level of accessibility, ease of retrieval, and overall cognitive salience (Harris, 1998; Bybee, 2010; Blumenthal-Dramé, 2012).

By considering frequency as an explanatory factor of linguistic behavior, usage-based linguistics has been very successful at offering comprehensive accounts of a range of thorny linguistic issues. Let us take the example of (ir)regularity in inflectional morphology. As already mentioned, the emergence of irregular forms in diachrony is accounted for by the fact that high-frequency forms are more resistant to change and may thus persist while lower-frequency forms are regularized. In addition, a usage-based approach can also account for the status of inflected forms in synchrony. Generative approaches traditionally treat regular and irregular inflected forms as qualitatively different: regular forms are derived by fully productive rules, while irregular forms are listed in the lexicon (e.g., Clahsen & Rothweiler, 1992; Pinker, 1991; Pinker & Prince, 1994). However, listing irregulars as exceptions fails to capture the fact that they may also follow sub-regularities, albeit of a limited scope, e.g., the English past tense forms *sing/sang*, *ring/rang*, *shrink/shrank*, *sink/sank*, etc. In fact, some of these patterns also exhibit limited productivity, especially with base forms that bear some similarity with a number of irregular items (Bybee & Slobin, 1982; Bybee & Moder, 1983; Albright & Hayes, 2003). For instance, a novel verb like *dize* might be assigned the past tense form *doze* instead of the expected “default” *dized*, by analogy with similar pairs like *rise/rose* and *stride/strode* (Albright & Hayes, 2003). As argued by Bybee (1995) and Baayen & Lieber (1991), the varying productivity of morphological patterns receives a usage-based explanation as a reflection of their type frequency, i.e., how many different base forms are attested in them (see also Wonnacott, Boyd, Thompson, & Goldberg, 2012). Bybee (1995) argues that high type frequency, in conjunction with phonological openness, leads to the emergence of a “default” pattern, such as the English past tense *-ed* suffix. Similarly, the co-existence of several high-type-frequency patterns leads to competition between them, with none standing out as markedly more productive than the others. Hence, the distinction between regular and irregular morphology in a usage-based account does not correspond to explicitly stipulated differences in grammatical status, but to quantitative differences that directly follow from usage and are matters of degree rather than categorical. In sum, a usage-based approach adequately explains the facts of inflectional morphology in a natural way, without recourse to arbitrary stipulation or innate categories.

As outlined in the previous section, the field of argument structure is also ripe with theoretical puzzles that have attracted a lot of attention in the literature. The purpose of this book is to investigate how a usage-based approach can address issues in argument realization and offer frequency-based explanations for its organizing principles. Many studies substantiate the claim that the structure of

grammar is ultimately tied to usage, and, presumably, no less may be said about the grammar of verbs. In fact, some of the findings mentioned above do point to the direct relevance of usage to argument realization (e.g., Gahl & Garnsey, 2004; Wiechmann, 2008; Wilson & Garnsey, 2009; Tily et al., 2009; Kuperman & Bresnan, 2012). Yet, many models still rely on traditional assumptions about the nature of syntax and lexis inherited from research prior to the usage-based turn, many of which have not been thoroughly examined from a usage-based perspective. This book seeks to mend this gap, by investigating usage-based principles of the organization of argument realization.

1.3 Structure of the book

The question of the usage basis of argument realization will be addressed at three levels of organization, from the more specific to the more general: first verbs, then grammatical constructions, and finally cross-constructional generalizations (traditionally known as argument structure alternations).

1.3.1 Part I: Verbs

All models of argument realization incorporate in one way or another the fact that verbs make reference to a number of participants. However, many verbs (if not most) can occur with more than one set of arguments (or valency pattern). As discussed in Chapter 2, so-called projectionist approaches claim that the verb is stored with several different argument structures, but in construction grammar, such an assumption is unnecessary, since, in that framework, constructions may select only a subset of the arguments of a verb for overt realization, or conversely, may contribute arguments by themselves. This, however, results into some indeterminacy as to how many arguments are contributed by the verb in a given instance, which in many cases cannot be decided on intuitive grounds, especially if it is assumed that a verb should only have one lexical entry. I suggest, following a proposal by Langacker (2009), that verbs can be stored with a range of valency patterns that is directly related to that verb's usage, in that repeated use of a verb in a construction leads to the conventionalization of the corresponding valency pattern with that verb, which as a result comes to be associated directly with it, and not arrived at via combination with a construction.

In Chapter 3, I test this hypothesis by comparing experimental results with usage data. I report on an incremental reading comprehension experiment designed to determine the relative cognitive salience of two valency patterns of three commercial transaction verbs: *buy*, *pay* and *sell*. I then compare the results of this

experiment to usage data extracted from corpora. As it turns out, the frequency of a verb in a specific syntactic environment correlates with the relative cognitive salience of the corresponding valency pattern for that verb (as measured by reading times), in that more frequent valency patterns are processed more easily. These findings provide evidence for the usage basis of verb valency.

1.3.2 Part II: Constructions

In Chapter 4, I turn to the second level of organization of argument realization: the generalization of argument structure over a number of verbs. In construction grammar, syntactic patterns of argument structure are paired with an abstract meaning. Previous research has shown that the meaning of a construction can be largely traced back to its verbal distribution (Goldberg et al., 2004; Stefanowitsch & Gries, 2003). This suggests that, from a usage-based perspective, a construction comes to be associated with the meaning of verbs most frequently occurring in it. However, I show that this view does not hold for constructions with a highly abstract meaning, such as the conative construction, which realizes the theme argument of a transitive verb as a prepositional phrase headed by *at* and is generally associated with an interpretation of attempted achievement of an intended result, as exemplified by such contrasts as *John kicked at the ball* vs. *John kicked the ball*.

In Chapter 5, it is argued that the distributional basis of such abstract patterns as the conative construction can be restored if we look at lower levels of generality. Argument structure can be simultaneously defined at several levels: highly abstract constructions (i.e., applying to a high number of verbs) at the highest levels of generality, and more specific ones (i.e., applying to a more limited number of verbs) at lower levels. Applying this principle to the conative construction, I show that if its instantiations by classes of semantically similar verbs are considered as independent constructions, the most frequent verbs of a given verb-class-specific construction are found to bear a close relation to the semantic contribution of the construction in that verb class. This means that it is possible to formulate constructional generalizations on the basis of verbal meaning at the level of clearly delimited verb classes, but sometimes not so easily at the most abstract level.

1.3.3 Part III: Alternations

The notion of alternations in linguistics refers to the possibility for the same unit to receive different formal realizations. In the grammar of verbs, pairs of constructions corresponding to different realizations of a common set of arguments have also been characterized as syntactic alternations, e.g., the dative alternation,

pairing the double-object variant (*John gave a book to Mary*) with the *to*-dative variant (*John gave a book to Mary*). In construction grammar, the variants of an alternation are seen as independent constructions with their own properties (cf. Michaelis & Ruppenhofer, 2001; Goldberg, 2002), but the relation between them is usually not considered part of the grammar. In Chapter 6, I question whether this view really provides a good account of what speakers know about their language. Adhering to Cappelle's (2006) allostructions model, I suggest that speakers store generalizations of a common constructional meaning over formally distinct argument structure constructions, and I report experimental evidence for this proposal. I also suggest that such alternation-based generalizations play a role in argument structure productivity.

In Chapter 7, I report on an experimental study testing for productivity asymmetries in the dative alternation and in the locative alternation (*spray paint on the wall* vs. *spray the wall with paint*). The experiment examines whether the extent to which speakers productively use a novel verb in the other variant of an alternation varies according to the variant with which this verb was previously presented to them. I find that there is indeed a productivity asymmetry in the dative alternation, whereby subjects are more likely to use the verb productively in the unattested variant when the novel verb was first presented in the double-object construction than when it was presented in the *to*-dative construction. In the locative alternation, no asymmetry is found. I argue that the productive behavior of speakers can be explained by patterns of usage. Drawing on corpus data, I show that many more dative verbs are attested in the *to*-dative variant only than in both constructions; conversely, almost all of the verbs attested in the double-object variant are also attested in the *to*-dative variant. In other words, the assumption that a given verb alternates is much more likely if that verb has been observed in the double-object variant than in the *to*-dative variant. Such patterns of type frequency are not found for the locative alternation, which explains its lack of asymmetry. I suggest that speakers have internalized these patterns of usage, and that their knowledge of grammar includes information about relations between constructions involved in an alternation.

In sum, the findings reported in this book show that the cognitive representations involved in argument realization are largely shaped by usage, at all three levels of organization. In Chapter 8, I elaborate on some theoretical and methodological implications of these findings.