

Chapter 1

Introduction

It is easy to take our knowledge of language for granted. We learn language before we carry our first backpack to school, and we use it almost every waking hour of every day. Although we may not have studied quantum theory, or read Homer or James Joyce, we are each expert at using our own native language. The challenge that all learners face becomes more apparent when we try to learn a second language in school or as adults.

There are many utterances that are perfectly understandable, but which nonetheless tend to be avoided by native speakers of English. If asked, speakers will agree that there is something mildly “off” about them, even though they may have difficulty articulating exactly why they don’t sound quite right. For example, we might confess that someone is driving us *crazy* (or *bananas* or *insane*), but we know that it would sound odd to complain that someone is driving us *angry*. We know that *tall* bushes are *high* bushes, but a *high* teenager is not necessarily *tall*. We can be creative in how language is used, but our creativity is constrained in ways that can be hard to articulate. For example, someone can *tell me something* or *tell something to me*, but they can only *explain this to me*; that is, it sounds somewhat unconventional to native speakers of English to say, *explain me this*. That is what this book aims to explain: when, why, and how native speakers are sometimes creative with language and yet at other times much more conservative.

Speakers avoid saying certain things, of course, simply because they want to avoid overtly negative reactions. The following are examples of such ill-advised utterances:

Sorry Mom, I didn’t mean to get caught.

I only care about my grade in this course.

Your nose is too big for your face.

But children are not systematically corrected for the types of utterances this volume aims to address, which will hereafter be indicated by a preceding “?” (*?explain me this*, *?drive him angry*, etc.). Caregivers are much more focused on the content of children’s speech than on its form, as long as the message is clear enough. For example, a child who says *Me loves you, mommy* is more likely to

get a hug than a grammar lesson, and a young child who utters an impressively grammatical utterance such as *I have just completed a mural on the living-room wall with indelible markers* is unlikely to get positive feedback from most parents. The sorts of formulations that native speakers recognize as odd are also not the sorts of formulations that grammar teachers warn against, since they are so rarely uttered by native speakers that no admonishment is needed.

To be clear, it is not that one *never* hears expressions such as *?explain me this* (or *?drive him angry*), or that all speakers judge them to be equally odd. In fact, speakers’ judgments are gradient and dependent on a number of interrelated factors that are the focus of this book. But corpus and experimental studies confirm that certain types of utterances are avoided by native speakers much more than would be expected by chance. In order to think about how these aspects of language are learned, it’s worth thinking about what speakers and language learners are trying to do.

1.1 The Puzzle

The learner’s goal is to comprehend messages, given the forms she witnesses, and to produce forms, given the messages she wants to convey. Therefore, speakers must learn the ways in which forms and functions are paired in the language(s) they speak. These learned pairings of forms and functions are referred to here as grammatical CONSTRUCTIONS. Speakers also aim to express their intended messages efficiently and effectively while respecting the conventions of their speech communities, as discussed more below.

Constructions generally allow us to apply our linguistic knowledge to new situations and experiences. English tends to be particularly flexible in the ways in which constructions are PRODUCTIVE. A few examples of productive uses of familiar constructions are provided in table 1.1, with labels for each grammatical construction provided on the right.

TABLE 1.1. Novel linguistic exemplars that demonstrate the productivity of various constructions

“Hey man, <i>bust me some fries</i> .”	Double-object construction
“Can we <i>vulture your table</i> ?”	Transitive causative construction
“Vernon <i>tweeted to say she doesn’t like us</i> .”	To infinitive construction
“What a <i>bodacious thing</i> to say.”	Attributive modification construction

Attested examples are cited in quotation marks. Here and below unless otherwise noted, attested examples come from Google.

TABLE 1.2. Novel formulations that are judged odd by native speakers

?She explained him the story. (cf. She told/guaranteed him the story.)	Double-object construction
?He vanished the rabbit. (cf. He hid/banished the rabbit.)	Transitive causative construction
?She considered to say something. (cf. She hoped/planned to say something.)	To infinitive construction
?The asleep boy (cf. The astute/sleepy boy)	Attributive modification construction

At the same time, the constructions exemplified in table 1.1 resist being used productively with certain verbs or adjectives, even when the intended meaning is perfectly clear. Examples that illustrate the lack of full productivity are provided in table 1.2. Under each ill-formed example is a closely related fully acceptable example, in parentheses. The latter are provided to indicate that there are no simple, system-wide explanations for why the odd sentences strike native speakers of English as odd. Thus, constructions can be extended for use with some words (table 1.1), but they are rarely completely productive (table 1.2), even when no general constraints are violated. How is it that native speakers know to avoid certain expressions while nonetheless using language in creative ways? It is no exaggeration to say that this basic question has bedeviled linguists and psychologists for the past four decades.

1.2 The Roadmap

The paradox of PARTIAL PRODUCTIVITY of constructions is what this book aims to address. We will also address several issues that have not widely been viewed as directly related. In particular, chapter 2 includes a discussion of how we learn to circumscribe the meanings of words. Close attention to word meanings reveals that speakers possess a vast amount of rich contextual knowledge about what each word means, and about which other words it tends to co-occur with. But, initially, young children make certain errors. They may call the moon a *ball*, or the mailman *Daddy*, before they learn and become fluent with other words (specifically, *moon* and *mailman*). That is, children need to learn to restrict their use of individual words by witnessing how those words *and other words* are used in particular contexts. The rest of the book argues that the same mechanisms involved in learning and restricting word meanings are used when learning and restricting grammatical constructions, and that

4 CHAPTER 1

this process explains how we come to avoid formulations such as *?explain me this*. By beginning with word meanings, I hope to make the discussion of our primary target—the partial productivity of grammatical constructions—more accessible. That is, once we have a better understanding of word meanings, we can tackle grammatical constructions by essentially asking: What would words do?

Chapter 3 outlines the various factors that are relevant to our knowledge of how grammatical constructions are used within a given speech community. These include formal properties (*syntax*), words and partially filled words (*morphology*), meaning (*semantics*), discourse function (*information structure*), and social context. An appreciation of these factors is a prerequisite for solving the *explain-me-this* puzzle. This chapter also highlights the remarkable degree of cross-linguistic variation that exists in how simple clauses are expressed in the world's languages, in an effort to emphasize just how much people must learn in order to use the constructions in their language appropriately.

The proposed solution to the partial productivity puzzle allows both generalizations (table 1.1) and exceptions (table 1.2) to be learned via the same mechanisms. In particular, in chapters 4 and 5, two key factors—*COVERAGE* and *COMPETITION*—are discussed. Chapter 4 explains how constraints on meaning and use *emerge*, as witnessed exemplars cluster within the high-dimensional conceptual space in which our representations for language exist. This chapter outlines how clustering licenses creative uses of constructions. In particular, a single factor, *COVERAGE*, combines variability, type frequency, and similarity; specifically, a new instance is licensed to the extent that the ad hoc category required to contain it has been well attested (has been sufficiently “covered”). Also outlined in this chapter is a useful model for formalizing the required mechanism; namely, an incremental Bayesian clustering algorithm (Barak et al., 2014, 2016; see also Alishahi and Stevenson, 2008; Matuszych et al., 2017).

In chapter 5, the critical role of competition is detailed. As we comprehend utterances, we attempt to anticipate what the speaker will say next, and we are able to use what the speaker actually says to improve future predictions through a process of error-driven learning. Repeatedly witnessing certain formulations in certain types of contexts strengthens the connections between those grammatical constructions and the intended messages-in-context expressed; this results in conventional formulations becoming more accessible for expressing the types of messages that have been previously witnessed. When there exists a readily available formulation that expresses the intended message in the given context, it usually wins out over potential novel formulations. A special effort is required to buck conventional formulations, although this is possible, for the sake of memorability or playfulness (as in the title of this book). But when there is no

readily accessible combination of constructions available to express a speaker's intended message-in-context, she needs to extend language creatively.

The proposal is situated in a larger context in chapter 6. Many studies have demonstrated that children are initially less creative than adults: children behave “conservatively” in that they generalize constructions less freely than adults do. Yet other studies have found that children generalize *more* broadly than adults. This apparent paradox is reconciled by recognizing that children are less adept at aligning bits of knowledge within their high-dimensional conceptual space: sometimes they fail to recognize relevant parallels across exemplars, at least with sufficient confidence (and so they behave conservatively); other times they fail to recognize or retain relevant distinctions (and so they generalize or simplify). Appropriate use of grammatical constructions emerges once the relevant conditioning factors for each construction are learned, and the language user becomes more fluent at accessing the appropriate constructions from memory.

Chapter 6 also outlines why adult learners of a second language tend to have particular difficulty avoiding the types of odd formulations this book addresses (including *?explain me this*). The suggested reasons go beyond the fact that adults receive less input overall, and that the input they do receive is less well suited to learning. In particular, adult learners need to inhibit their well-practiced native language in order to process a new language, and this appears to lead to a reduced ability to take full advantage of the competition among constructions within the new language. Since competition is argued to be key to constraining generalizations via statistical preemption (chapter 5), second-language learners tend to be more vulnerable to producing certain types of formulations that make sense but which native speakers systematically avoid. Additionally, while adults are generally quicker to discern which dimensions of similarity and dissimilarity are relevant to clustering linguistic representations within their hyper-dimensional conceptual space, they are at the same time prone to miss very subtle similarities and distinctions that are not relevant in their first language.

1.3 The CENCE ME Principles

The basic understanding of language that this book outlines is based on the key ideas listed in table 1.3, which are discussed in detail in the following chapters. An acronym of the key words in these principles is EEMCNCE, but EEMCNCE would be impossible to pronounce. So, let us instead use an anagram of EEMCNCE: **CENCE ME**. “CENCE ME,” pronounced “sense me,” is intended to emphasize the importance of sensible communication. CENCE ME also usefully illustrates productivity, since the phrase itself is a novel use of the transitive construction. The CENCE ME principles spell out some key assumptions of the more general

TABLE 1.3. The CENCE ME principles

-
- A. Speakers balance the need to be Expressive and Efficient while conforming to the conventions of their speech communities.
 - B. Our Memory is vast but imperfect: memory traces are retained but partially abstract (“lossy”).^a
 - C. Lossy memories are aligned when they share relevant aspects of form and function, resulting in overlapping, emergent clusters of representations: Constructions.
 - D. New information is related to old information, resulting in a rich network of constructions.
 - E. During production, multiple constructions are activated and Compete with one another to express our intended message.
 - F. During comprehension, mismatches between what is expected and what is witnessed fine-tune our network of learned constructions via Error-driven learning.

^a Our representations are “lossy,” a term from computer science, in the sense that they are not fully specified in all detail.

USAGE-BASED CONSTRUCTIONIST APPROACH to language that are widely shared (see, e.g., Bybee, 2010; Christiansen and Chater, 2016; Goldberg, 2006; Kapatsinski, 2018; Langacker, 1988; Tomasello, 2003; Traugott and Trousdale, 2013). The approach also shares much with memory-based EXEMPLAR-BASED MODELS (Aha et al., 1991; Bod, 2009; Bybee, 2002; Daelemans and van den Bosch, 2005; Gahl and Yu, 2006; Kruschke, 1992; Nosofsky, 1986). The CENCE ME approach emphasizes that exemplars—structured representations—cluster within a hyper-dimensional conceptual space giving rise to emergent constructions, which are then extendable as needed for the purpose of communication.

Individual languages can and do vary in striking ways, as will be emphasized, but the usage-based constructionist approach adopted here suggests that the CENCE ME principles are at work in every natural language, serving to constrain and shape the range of possible human languages. The present book emphasizes examples in English because the majority of the experimental and modeling work to be described has been done on English, and because English is the language I know best.

My understanding of what a construction is has evolved. Early on, I adopted the following definition:

C is a CONSTRUCTION if and only if C is a form-meaning pair $\langle F_i, S_i \rangle$ 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)

Later, I recognized that this definition was too narrow. Our knowledge of language comprises a network of constructions, and we clearly know and remember conventional expressions even if they are in no way idiosyncratic. So I broadened my definition of constructions as follows:

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).

The present volume offers a still more inclusive understanding of what constructions are, motivated by a better appreciation of human memory, learning, and categorization. Here, as explained in the following chapters, constructions 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.

Proponents of alternative perspectives or readers who wish to compare the present proposal with other proposals in more detail may find chapter 7 particularly relevant. There, several recent alternative proposals that aim to account for the partial productivity of constructions are discussed. These include, for example, the idea that speakers avoid straying from what they have witnessed ("conservatism via entrenchment"), that it is useful to posit invisible syntactic diacritics or underlying structures without specifying how these are to be identified by learners, that putting a cap on the number of exceptions and a floor on the number of instances that follow a generalization will ensure how and when generalizations are productive (the Tolerance and Sufficiency principles of Yang [2016]), or that incorporating degrees of uncertainty into formal rules is predictive (O'Donnell, 2015). While aspects of each of these proposals have merit, we will see that the usage-based constructionist approach, described by the CENCE ME principles, explains the facts more fully. The final chapter stands back and puts the discussion in a broader context, while raising several outstanding issues that remain to be addressed.

1.4 Speakers Are Efficient and Expressive and also Conform

Before leaving this introductory chapter, let's go over the first of the CENCE ME principles: We aim to express our messages effectively and efficiently while

8 CHAPTER 1

obeying the conventions of our speech communities. To clarify the key terms involved:

1. **Expressiveness:** Linguistic options must be sufficient for conveying speaker's thoughts, beliefs, and attitudes in ways that listeners are able to understand.
2. **Efficiency:** Fewer and shorter constructions are easier to learn and produce than more or longer constructions.
3. **Obeying conventions:** Learners attempt to use language in the ways that others in their language communities do.

A language is only sufficiently expressive if it has the means to adequately convey a speaker's thoughts, beliefs, and attitudes in ways that avoid failures of communication. A *maximally expressive* language might have an ever-increasing number of words and constructions, with every potential distinction indicated by a unique form. On the other hand, a *maximally efficient* language would have a single, easy to learn and use form (perhaps the form, *ah*). The fact that language users need to be both effective *and* efficient requires natural languages to find a balance between these two opposing factors, as has been long discussed by functional linguists (Briscoe, 1998; Bybee, 1985, 2003; Givón, 1979; Goldberg, 1995; Grice, 1975; Haiman, 1985; Levinson, 1983; Paul, 1888; Slobin, 1977; von Humboldt, [1832] 1999).

The recognition that languages must be efficient and expressive, and that these pressures are mutually constraining, has gained new traction within the “noisy-channel” approach to language processing, which recognizes that speakers are attempting to express information as efficiently as possible, under imperfect or noisy conditions (Gibson et al., 2013; Jaeger and Levy, 2006). The noisy-channel approach has emphasized the dynamic nature of the balance between efficient and expressive communication. When a speaker is reasonably certain that an intended message will be successfully conveyed, the balance is tipped toward efficiency, with forms being reduced and distinctions being underspecified (Hopper and Traugott, 2003; Jaeger, 2010; Levy, 2008; Levy et al., 2009; Piantadosi et al., 2011, 2012). For example, when a verb appears with the construction it most commonly appears in, the verb form itself tends to be reduced (Gahl and Garnsey, 2004). Similarly, the complementizer *that* is more likely to be omitted when it is predictable in context (Wasow et al., 2011). On the other hand, under noisy or uncertain conditions, distinctions may be exaggerated, and language may be made less ambiguous in various ways (Bradlow and Bent, 2002; Buz et al., 2014; Gibson et al., 2013; Hall et al., 2013). Thus, efficiency and expressiveness balance each other and lead languages to vacillate between using shorter or fewer forms to express a given message on the one hand, and

adopting new, longer, or additional forms in order to ensure that messages are understood as intended, on the other.

The last idea, that speakers tend to obey the conventions of their language community, captures the fact that humans treat language as a normative enterprise. For example, people within a given community tend to believe that there is a “right” way to pronounce words, even if other communities are recognized to pronounce them differently. The Gershwins’ famous lyric, “You like tomato and I like tomahto, . . . Let’s call the whole thing off” epitomizes this idea. Humans are a rarity within the animal kingdom in using arbitrary communicative symbols that are shared within a community and distinct from those used in other communities (Tomasello, 2016). In fact, humans quite generally appreciate that there are culture-specific “right” and “wrong” ways to do a great many things, and we learn to obey these conventions in a way that other species do not (Boyd and Richerson, 1988; Horner and Whiten, 2005). Many normative conventions are, at least initially, self-conscious; for example, our knowledge about how to eat food politely, whether or how much to tip at restaurants, or whether it is polite to sneeze or burp in public. Other social norms may be obeyed without self-conscious awareness, including how close to stand to each other while speaking, or what sort of foods are appropriate for breakfast.

The importance of cultural norms for human behavior has enjoyed a long and rich appreciation within philosophy (e.g., Korsgaard et al., 1996). Our respect for normative patterns of behavior is what allows us to create complex cultural practices. For example, dollar bills would be meaningless were it not for the social agreement that imbues them with value. Driving would be terrifying if we couldn’t rely on other drivers to (generally) obey standard driving practices. Again, work that has compared humans with other primates has emphasized that cultural norms may be uniquely human (Tomasello, 2009, 2016), particularly when they serve no clear purpose (e.g., Horner and Whiten, 2005; McGuigan et al., 2011).

Normativity is critical to the *explain-me-this* problem in that generations of learners obey restrictions that do not serve any clear communicative function: we respect the patterns that are evident in the input. For example, even though saying *she made it vanish* is somewhat less efficient than *she vanished it*, and even though the latter formulation is readily interpretable, native English speakers avoid the shortcut and normatively obey a shared preference for the periphrastic form. A mechanistic explanation supports our tendency to produce forms that our community deems acceptable, insofar as partially familiar formulations are easier to access than wholly novel formulations. That is, more conventional forms are more efficient to access from memory even if they are less efficient to produce (see section 4.5). But the reason native speakers judge the longer phrase (*explain this to me*) to be the “right” way to express the intended meaning

10 CHAPTER 1

and the shorter phrase (?*explain me this*) to be “incorrect,” and the reason that familiarity tends to be more important than ease of articulation, is because we desire to speak like others in our community—language is a social and normative enterprise.

Chapter 2 briefly reviews the nature of word meanings and asks how we learn to use words appropriately. We will then see in the following chapters that many of the lessons learned from an appreciation of word meaning extend naturally to our primary question: How do we learn to use basic clause types—ARGUMENT STRUCTURE CONSTRUCTIONS (ASCs)—in creative but constrained ways?