Doctoral Thesis

On the Cognition of the Expression of Thought in Language

Ruben Bernard Polo-Sherk

Division of Integrated Arts and Sciences
Graduate School of Integrated Arts and Sciences
Hiroshima University

March 2021

Contents

Acknowledgements

I am grateful to my advisor, Akira Machida, for his guidance during my studies and for providing a stimulating and comfortable environment that enabled me to conduct my research fruitfully. My experience was a thoroughly enjoyable and fulfilling one.

Also thanks to my sub-advisors, Shinpei Tanaka and Miki Shibata, and the other three members of my committee, Robert Taferner, Shogo Sakata, and Makoto Iwanaga, for discussion and help with proofreading. Their comments improved the manuscript in several places.

I benefited from engaging discussion with Russell Kabir and Aaron Sponseller on my ideas at an earlier stage. Noriko Yamane and Jon Clenton also provided helpful input.

Finally, I would like to thank my family who helped with the making of this work. Discussions with Zheng Xinshuang helped clarify and develop some points of my theory and solidify how to explain it, and she also checked my Chinese examples. Peter Sorek checked my Hebrew examples. Annabel Sherk and Juan Polo let me talk out what was in my head, which helped me figure out how to explain my ideas and find suitable wording, provided feedback on understandability on what I wrote, and proofread the manuscript.

List of glossing abbreviations

1	first-person
3	third-person
ABL	ablative
ABS	absolutive
ACC	accusative
ADJ	adjective
ADV	adverb
ANT	anterior
COMP	complement
COP	copula
CLF	classifier
DAT	dative
ERG	ergative
F	feminine
GEN	genitive
IMP	imperative
LOC	locative
M	masculine
NEG	negative
NMLZ	nominalizer
NOM	nominative
NPST	non-past
OBJ	object
PFV	perfective
PL	plural
POSS	possessive
PRF	perfect
PRS	present

PST	past
Q	question marker
SBJ	subject
SG	singular
SIM1	simultaneous 1 (continued simultaneity)
SIM2	simultaneous 2 (interrupted simultaneity)
TOP	topic

List of Figures

List of Tables

Chapter 1

Introduction and Overview

1.1 Introduction

Understanding the process of the expression of thought in language is of central importance to our understanding of language, thought, and cognition. This issue relates to the function of language as a system, cognitive processes, and how they interact. The basic question to be answered is: What are the cognitive processes involved in creating a linguistic expression from thought and how do they achieve it?

Since thought is not limited by language, the general case we must deal with is the derivation of linguistic expressions from language-independent thought. That is: the original thought is not in the target language and must be formulated in it; our goal is to describe the cognition responsible for deriving the linguistic expression from the thought. Our target is not what the substrate of thought (i.e. the 'language of thought') is itself, but the cognition of the 'translation' from thought to language.

Fundamentally, linguistic expressions are arrangements of words, morphemes, and constructions that yield particular meanings; the task of expressing thought in language, then, can be described as coming up with an arrangement of these elements that will yield an appropriate meaning. Our approach to investigating this will therefore involve exploring meaning construction in language (the mechanisms by which arrangements of linguistic elements yield particular meanings), explicating what is involved in constructing arrangements of linguistic elements, and identifying the cognition used to find a suitable arrangement. We will see over the course of the discussion that this cognition is a kind of abduction (the production of a cause for a desired target effect) that is driven by background knowledge of meaning construction.

The central piece of this dissertation is that abduction is the primary cognitive operation that effects the expression of thought. I explore the details of how this cognitive operation functions and develop a model of the process of expressing thought in language.

1.2 Overview

Chapter 2

First, however, we must properly set up how the problem should be characterized and what our approach to it should be. There are a number of background issues that must be explored to do so.

Section 2.1

The starting point for our inquiry is the independence of thought and language—more precisely, the issue of whether thinking is done in some language or is originally language-independent. I first discuss various views on the (in)dependence of thought and language. It is a common intuition that thinking is done in language, but this is not actually the case, at least in general. In essence, while language can influence thought (e.g. by providing patterns that can serve as templates for thought and by being a vehicle for concept acquisition), thought itself is in general independent of language. This gives rise to the problem that is the subject of this dissertation: since thought is originally not formulated in language, linguistic expressions must be derived somehow from the original thought. At the same time, this discussion brings up some effects that represent influence from language on thought and whose consideration is necessary for a comprehensive description of the cognition of expressing thought in language.

Section 2.2

The next topic to be addressed is the nature of linguistic representations—what do expressions mean? Making this clear is important for characterizing the content of linguistic expressions, and, ultimately, its relation to thought. Several issues are

examined. First is the question of whether grammatical elements and lexical elements are fundamentally different and should be treated in isolation. Whether the lexicon and grammar should be treated with distinction in theoretical description is a point of some debate; I discuss this in the context of our problem, and argue that they should be treated in an integrated manner. The second is the question of what counts as a linguistic element. Related to the previous point, elements include words, morphemes and constructions. But there are complicating factors that make this not quite so straightforward a definition. The short version of the answer to this question is that what counts as an element and what meaning it has depends on the language user's analysis. Next is the granularity of the conceptual structure expressed in language (i.e. how precise it is) and how this compares to that of thought. On one hand, linguistic forms are often underdetermining with respect to their referents, but they can also be more specific than the thoughts they express. Finally, I address the issue of what the content of linguistic meaning is; I discuss four important aspects: propositions, imagery, metaphor, and structural relationships.

Following this, we consider the problem of the format of the mental representation of concepts. This is a complex and unresolved matter. A major debate is whether concepts are mentally represented with perceptual information or not. There has been a great deal of investigation concerning this in cognitive neuroscience and cognitive psychology. I discuss various problems regarding the hypothesis of a perceptual format of conceptual representation, and conclude that while it is possible that perceptual information may have some role in the mental representation of concepts, it probably cannot be primarily constitutive.

The last component I address here is lexicalization (the process of encoding ideas in the lexicon). The cognitive ability underlying lexicalization is what I call *conceptual lexicalization*—the ability to take any bit of thought and index it as a concept (without necessarily assigning it to a linguistic element). This has important implications for how language-independent thought can be organized and related to linguistic structure.

Section 2.3

We must also address the issue of the architecture of language—the question of how the 'grand structure' of language (essentially, how linguistic structure relates to meaning) should be characterized. Making clear what sort of architecture is suitable for addressing our problem sets up how the problem should be seen, and therefore how it should be approached. The discussion here draws on many of the issues discussed earlier. I review various conceptions of the architecture of language, discussing questions of how linguistic form relates to meaning, and clarify the architecture that will be adopted for this dissertation.

Section 2.4

All of this takes us to the final piece of the background: what the basic schema of the process of expressing thought in language is. This will clarify the pieces we need to examine and account for in order to answer the question of what the cognition of expressing thought in language is. I discuss ideas from previous research that are related to this issue. I then outline the process as I see it. The basic process argued for is as follows: The language user starts with a language-independent thought and must make an arrangement of linguistic elements that generate an approximation of that thought. A 'prerequisite' to making an arrangement of elements is organizing the information to be expressed in accordance with the mechanisms of meaning construction; this organization allows selecting the linguistic elements that are to comprise the arrangement (lexical selection) and structuring those elements. The components to this process will be explored in detail in later chapters; establishing what the process is allows us to finally directly address the question of what cognition is used to carry it out.

Chapter 3

The first piece of the problem we take up on the way to the solution is meaning construction. The essential problem we must address is the (obvious) fact that particular arrangements of linguistic elements generate particular meanings. This is not a straightforward matter. Composition of meaning does not generally result in a meaning

comprised of a list of the meanings of the components, and meaning is sensitive to linguistic context: that is to say that meaning construction is dynamic. I discuss the issue of compositionality, and then address the problem of the cognitive mechanism of meaning construction. I review theories on the cognition that forms the mechanism for meaning construction, whose proposals include inference and simulation. I conclude with a proposal that the cognition driving the collective dynamic rules of meaning construction is based on the understanding of how linguistic meanings function and interact; this cognition is related to that used for making sense of mental models, and can be connected as well to both inference and simulation.

Chapter 4

We then explore the details of constructing an arrangement of linguistic elements. In a most basic sense, this consists of organizing the information to be expressed and selecting the linguistic elements that are to constitute the arrangement. The organizational structurings of semantic material contained by linguistic expressions reflect the mental organization of conceptual information as expressed in language, of which there is a wide variety within languages and especially across languages. I review research that addresses aspects of this issue of mental organization, and then introduce *compartmentalization*, a general theory of the mental organization of information in language in accordance with meaning construction. For our purposes, an understanding of the organization of information that fails to take into account meaning construction is insufficient: information is necessarily organized such that a suitable meaning is constructed. Compartmentalization describes both the structuring of linguistic expressions as well as part of the process of expressing thought in language. The selection of linguistic elements is a complex issue. I first summarize our understanding of lexical access, including theories of associative access (spreading of activation), and discuss priming effects. I then discuss categorization, the consensus account of lexical selection in cognitive linguistics. Finally, I argue that the selection of linguistic elements for expressing thought should be understood through compartmentalization, rather than categorization alone.

Chapter 5

This puts us in a position to consider the cognition used to create a linguistic expression from thought. The central claim of this dissertation—the answer to our original question—is that the principal cognitive process responsible for the expression of thought in language is a kind of abduction. Abduction is a form of inference that involves reasoning from effects to causes. I first give a basic introduction to what abduction is and what an abductive problem is. I then clarify abduction as a matter of logic versus abduction as a matter of cognition: we are interested in the latter, which is the production of a cause that would yield some desired effect. Particularly in the context of logic, abduction is sometimes referred to as "explanation", "hypothesis generation", or "inference to the best explanation". I explain that all of these miss the mark for describing the cognitive operation of abduction. Many cases of abduction—including abduction for the expression of thought in language—have nothing to do with explanation. Hypotheticality in inference is resultant from uncertainties in the logical system and is not a characteristic of abduction itself—deduction and induction can just as well generate hypothetical conclusions. Additionally, as a cognitive matter, the conclusion produced by an abduction is not a claim about the truth of a proposition, but rather a claim that some cause would yield some effect. Inference to the best explanation involves evaluating how good of an explanation various hypotheses are (after generating those hypotheses); all three of these aspects—explanation, production of hypotheses, and retrospective evaluation of them—are not involved in the actual cognitive operation of abduction. Following this, I show that expressing thought in language is abduction and clarify what sort of abduction we are interested in. Abduction for the expression of thought in language involves a 'cause' and 'effect' that are approximately the same thing (which is trivial in terms of formal logic, but quite significant when dealing with transformations of representations), the creative production of the solution (rather than its selection from a list), and reference to a background theory to determine what a potential cause would yield (rather than reference to pre-established rules specifying particular effects for particular causes). Next, I explore the mechanics of this abductive cognition, addressing five main aspects.

Abduction is closely connected to the psychological phenomenon of insight, though not wholly identical with it. Within this discussion, I clarify the relationship between abduction and insight and refer to research on insight to give a further-enriched understanding of some of these aspects of the mechanics of the cognition. First, I propose that the particular kind of abduction used in expressing thought involves creative manipulation of mental models. Next, I point out that abductive problems can be solved in one step or in multiple steps. Thirdly, I discuss the mechanism for searching for elements with which to build the solution/linguistic expression. Fourthly, I suggest a few cognitive mechanisms that may account for automatizing the process of expressing thought in language. Lastly, I consider the potential advantage that unconscious processing may have for avoiding working memory capacity limits, which may otherwise be a hindrance. In the final section of the chapter, I describe a model of the process of expressing thought in language.

Chapter 6

In the concluding chapter, I summarize the investigation. I then speculate on some important issues that are beyond the direct scope of the dissertation, but nevertheless are related to ideas brought up in the dissertation: the importance of conceptual lexicalization in thought in general, and the much-debated question of what comprises the language faculty. I conclude by noting directions for future research.

Chapter 2

Background

2.1 The (in)dependence of thought and language

The first question we must address is whether thought as originally conceived is formulated in language or in some language-independent form. If thought is dependent on language, there is no problem of how it is linguistically expressed, since from the outset it is conceived of in language. If, however, thought is free and independent of language, a problem arises: How is a linguistic expression constructed based on that language-independent thought?—i.e. the subject of this dissertation.

The fundamental predication of the theory in this dissertation is that thought is language-independent. However, a number of philosophers have argued that language is the medium of thought. First (Section 2.1.1), we will review their position. Next (Section 2.1.2), we will discuss reasons to believe that thought is indeed language-independent. Then (Section 2.1.3) we will note some caveats to the language-independent stance—ways in which language may actually influence thought. Following this (Section 2.1.4), we will discuss the relationship between conceptual structure and linguistic expressions: expressions depict conceptual structure, and this conceptual structure will generally be quite similar to—though not quite the same as—that of the original thought. The last subsection (Section 2.1.5) summarizes the preceding discussion and lays out the position that I adopt.

First, a few points to set up the discussion:

To begin with, we must specify the sort of thought we are concerned with. We are concerned with what I will refer to as *linguistic thought*. Linguistic thought is (without pre-judging the issue) thought of the sort that can be expressed in language. Now, there are many kinds of thought that cannot possibly be conceived of in a linguistic

formulation: e.g. the thought behind a musical composition or a painting cannot be expressed in language (and most of it cannot even be 'explained' in language[1]); these are, of course, irrelevant. But at the same time, clearly there are also thoughts that can be expressed in language. So, we can say, at the very least, that thought expressed by language is a certain 'type' of thought: this is linguistic thought.

Many researchers refer to what I call 'linguistic thought' as *propositional thought*. I do not adopt that description, because it is questionable whether propositionality characterizes linguistic thought. I am doubtful, at least, of its validity as a complete description: for instance, individual lexical concepts are certainly 'linguistic' (e.g. the meaning of *beauty* is linguistic in a way that the musical 'meaning' of a particular chord in a particular spot in a piece of music is not), but not necessarily propositional. In Sections 2.1.4 and 2.2.1.4 in particular, we see several kinds of linguistic meaning or aspects of linguistic meaning that are not really 'propositional'; it would seem to be a mistake to assume that the underlying thought is appropriately (or at least exhaustively) described as propositional. Now, one can say that linguistic meanings have to be of a type that is suitable to *participate* in propositions. But then propositionality is quite unhelpful as the defining characteristic of either linguistic meaning or linguistic thought, since 'propositional' would effectively mean 'linguistic'; the definition is circular, and 'propositionality' adds nothing. (See Section 2.2.1.4 for more discussion of propositionality.) At any rate, for the discussion in this section, the distinction between 'propositional thought' and 'linguistic thought' is not important. 'Propositional thought' referred to in the reviews below should be taken to be interchangeable with 'linguistic thought'.

The second point is a clarification of what is meant by 'language-independent'. What I present in the dissertation as a whole does not depend on the assumption of a particular format of thought. There have been a variety of proposals concerning the format of linguistic thought: it has been described as being comprised of symbols

[1] But see the discussion in Section 2.1.4.2 (Non-linguistic conceptual structure).

forming an abstract language[2] (e.g. Fodor 1975), as being comprised of perceptual symbols (e.g. Barsalou 1999), as abstract iconic representations of states of affairs (e.g. Johnson-Laird 1983), and as a non-propositional "unique whole" (Croft 2007). It is possible to draw a distinction, as Machery (2005) does, between *sententialist* views, which hold either that thought is conducted in language or that it is comprised of elements that form an abstract language, and all other views (Table 1). For our purposes, however, it is sufficient that the original thought be different from the mental representation of the ultimate linguistic expression; this is what I mean by 'language-independent'. I do not require that it be non-sentential—only that it be not in language (Table 2). In this dissertation, I will touch on issues relating to thought itself, but giving a characterization of thought is not the objective, and is not necessary for saying what I want to say.

Table 1: Machery's (2005) treatment of views on the format of thought

Sententialist views	Thought is in language
	Thought is in an abstract language
Non-sententialist views	Thought has some other form

Table 2: The treatment of views on the format of thought in this dissertation

Thought is dependent on language	Thought is in language
Thought is independent of language	Thought is in an abstract language
	Thought has some other form

[2] An abstract language is a symbolic system formed from a set of fixed meaningful atomic symbols and fixed means of composing them; the meaning of an 'expression' in such a system is determined by the symbols and how they are combined. Crucially, the symbols and means of composition are not necessarily the same as those of actual languages.

Finally, it should be noted that the peripheral issue of the relation of thought and language to consciousness figures quite importantly in many of the ideas reviewed in the following discussion. I do comment on this, but the overall theory in this dissertation is not contingent on whether the language user is consciously aware or unaware of the original thought or its linguistic expression.

2.1.1 Thought is dependent on language

Carruthers has offered a few variations of the idea that thought is conducted in language. In his 1996 book, he argues that thought (particularly conscious thought) uses language as its medium. The principle basis for this is introspection: people are generally aware of their conscious thoughts as having a linguistic form; therefore it can be assumed that that linguistic form constitutes those thoughts. As he states in the introduction (p. 2), "When a speaker utters a sentence...their utterance expresses a thought by *constituting* it, not by encoding or signaling it." Carruthers also links this linguistic medium to the establishment of consciousness, proposing that it does so by enabling thought that references other bits of thought[3].

In his 2002 article, Carruthers advances the idea (based on the assumption of a modular mind[4]) that language is the medium for not just conscious propositional thought[5], but for all thought involving integration of information from multiple modules. Intra-modular concepts are represented in a non-linguistic format and are generally

[3] See Clark (1998) for a similar idea on the relation of language to "thinking about thinking". Also see the discussion on *conceptual lexicalization* below (Section 2.2.3); one of the points elaborated on there is that reference in thought to cognitive objects does not require linguistic encoding.

[4] Modularity of mind, first proposed by Fodor (1983) is the hypothesis that the mind is partitioned into a number of autonomous 'modules' that are each responsible for a particular function and process different sorts of information.

[5] For Carruthers, propositional thoughts are "conceptualized propositional states, such as beliefs, judgments, and desires"; these are to be distinguished from mental model representations, percepts, and other imagistic representations, because mental models, percepts, and images by themselves are not propositional (2002: 2, 63 note 3).

independent of language (see 2002: 62 note 7); integrating concepts from multiple modules requires translating the representations of the concepts into a cross-modularly-compatible format—into a linguistic expression. Now, non-linguistically-encoded concepts—and thoughts themselves (as long as they are not cross-modular)—can be non-linguistic: within a particular module, both conceptual representation and generation of thoughts are possible. So, while not all propositional thought is conceived of in language, all cross-modular propositional thought is. Carruthers (2002: 13) gives the example of a thought to the effect of THE TOY IS TO THE LEFT OF THE BLUE WALL; this thought can only be conceived of in the form of a linguistic expression, because it contains both geometrical information and color information, and thus is cross-modular. The geometry module is capable of thoughts like THE TOY IS IN THE CORNER WITH A LONG WALL ON THE LEFT AND A SHORT WALL ON THE RIGHT, and the object-property module can create thoughts like THE TOY IS BY THE BLUE WALL; these thoughts are intra-modular and do not have to be represented in language. In order to conceive of a thought like THE TOY IS TO THE LEFT OF THE BLUE WALL, however, the intra-modular component thoughts must be inputted into the language system, which will integrate them and produce a thought in a cross-modular format. Carruthers proposes that the format in which this integrated thought is represented is the logical form (in the conception of Chomsky 1995) of the expression. One role that syntax would play in this sort of integration is in coordinating the indexing of any conceptual objects that are shared across modules (e.g. in the above example, WALL appears in both the geometrical thought and object-property thought). In Carruthers' account, all cross-modular thought is conducted in language; this thought can be conscious or unconscious. The logical form representation of a thought would include lexical items and syntax, but no phonological features. When the logical form is given a phonological representation, the thought will be conscious[6]. The specific mechanism for conscious awareness that Carruthers proposes is that the phonological output of the language

[6] This point is quite similar to Jackendoff's conception (discussed in Section 2.1.2), though Carruthers' and Jackendoff's positions on the (in)dependence of thought and language are fundamentally different.

system can be accepted as input by the comprehension sub-system, and thereby consciously perceived.[7]

Frankish (2002) generally agrees with Carruthers' proposal, but denies the existence of non-conscious propositional thought: no thought is represented as logical form alone without phonology; all cross-modular thought is conducted in language, represented phonologically, and conscious. In his 1998 article, Frankish provides a more indirect argument. He makes a distinction between cognition involving only "low-level" belief and "deliberate" *premising* (from Cohen 1992), which is contingent on "actively-formed" *virtual beliefs*[8]. The former is allowed to be non-linguistic, but the latter is argued to involve language. The key bit of reasoning is that premising requires ("implicit" or "procedural") inference rules, which are held by Frankish to be effected by linguistic knowledge; the premises themselves also are represented linguistically, since in order to have inference rules applied to them, their formal properties must be accessible, and it is presumed that this means that they are represented linguistically. Frankish (2018) identifies inner speech as conscious propositional thought, and links

[7] Numerous criticisms on various aspects of Carruthers' proposal can be raised (see, for example, the many commentary articles in that journal issue). The major criticisms that are relevant to our inquiry are:

> i) In looking for an answer to the problem of cross-modular integration, Carruthers skips all the way to language. Bickerton (2002) points out that an abstract language of thought can fill the role that Carruthers posits for linguistic expressions, and can do so in mostly the same way that Carruthers imagines language does. Molina (2002) notes that there is no reason to discount a pre-verbal level of thought (i.e. what I call the *original language-independent thought*) and that Carruthers provides no justification for the idea that thought must be conducted in linguistic expressions.
>
> ii) Slobin (2002) points out that different languages differ on the specifics of how a particular effective meaning is represented; Carruthers' hypothesis entails that thought be determined by the particularities of the language.

[8] On the distinction, Frankish says: "The principal difference between them will be that the premiser deliberately *guides* their inferential processes in order to keep them in line with the premise, whereas the ordinary believer leaves them to subpersonal control" (p. 7).

use of language in problem solving to Type-2 processes[9]. He employs (obliquely) a version of the argument from introspection. He argues that conscious thought (and Type-2 processing) requires a symbolic medium that is perceptually or introspectively available, and that language provides this.

Gauker (1994, 2002, 2011, 2018) has likewise contended that language is the medium of conceptual thought. A major part of his theory is that language is necessary for discrete conceptual categorization, and therefore for cognition involving discrete conceptual categorizations, which he distinguishes from imagistic thought (which does not involve discrete categorization)[10].

2.1.2 Thought is independent of language

We will next review the position that thought is independent of language—that is: that thought is not conducted in language.

The most fundamental problem is that the position that thought is conducted in language restricts thought to manipulations of whatever concepts[11] exist in whatever language is held to be that which thought is conducted in: if it is not in that language, it cannot be thought. Of course, however, people are free to think whatever they like, and are not rigidly or deterministically restricted in this manner. The experience of having a

[9] *Dual-process* theories of reasoning distinguish between *Type 1* processing, which is fast, automatic, and unconscious, and *Type 2* processing, which is slow, controlled, and conscious. See Frankish and Evans (2009) and Kahneman (2011) for overviews.

[10] Many authors are unconvinced of this assertion (see Carruthers' response (2002: 57) to Gauker 2002, and Machery's (2012) and Briscoe's (2014) reviews of Gauker 2011). Again, see the section on conceptual lexicalization (2.2.3) for a discussion on non-linguistically-encoded conceptual representations.

[11] In language, meaning is not encoded only in word-level units; morphemes and constructions also encode meaning (this why I refer to *linguistic elements* as words, morphemes, and constructions in this dissertation—see Sections 2.2.1.1 and 2.2.1.2). The commonly-used term *lexical concept* can be misleading in this regard (though it is often taken to include concepts associated with morphemes and fixed phrasal expressions); since I am referring here to all concepts/meanings available, I intend to include non-lexically-encoded concepts.

thought clearly defined in one's mind but not being able to find a suitable expression is almost certainly universal—this alone is fatal to the notion that thought is conceived of in language. There is evidence that pre-verbal infants have conceptual systems that can be used for categorization and problem-solving (see Mandler 2004 for a review) (although 'pre-verbal' infants may have acquired language to a degree (Friederici et al. 2011 exposed 4-month old German infants to grammatical and ungrammatical patterns in Italian and found they could distinguish between them after 'training')). Casasanto (2016) reports findings that people often use, e.g., left-right mental timelines (as seen in gesture) that have no corresponding metaphorical instantiation in their language. Levinson notes that languages often clearly lack lexical concepts that other languages have: for example, there may be no separate terms for 'blue' and 'green'—instead, the concepts will be covered by a single term 'grue'; there are varying degrees in lexical precision for kinship relationships (for instance, some languages have a term for 'father's sister's husband', and others do not); some languages have explicit encoding of 'and', 'if', and 'or', and some do not; but lack of lexical coverage in one's language does not mean that those concepts cannot feature in a person's thought (2003: 292-293). Nuyts (2012) brings up the fact that the expressive capacity of language is often deficient; he gives the example of the taste of garlic, jalapeños, and cilantro, which can only be described clumsily and roughly; they can, however, be conceptualized quite cleanly, of course: so thought involving the specific perceptive qualities of these objects cannot possibly be conducted in the medium of language (p. 330).[12] As Levinson states

[12] Nuyts (2012: 224-228) also makes an argument for the independence of thought from language based on the example of intra-language diversity in epistemic modality expressions. The idea is that various 'alternative' expressions (like *I <u>believe</u> he went to the bakery* / *It is quite <u>likely</u> that he went to the bakery* / *He <u>probably</u> went to the bakery*) encode the same central epistemic notion and differ only in (inter)subjectivity, performativity, and pragmatic information structure; they encode the central epistemic modality and the peripheral pragmatic factors all together and without reference to other alternative expressions, as "lexically/grammatically basic" expressions: the original conception of epistemic modality, therefore, cannot have the format of any specific expression. This argument, however, rests on two assumptions that probably should not be made, at least without qualification:

(2003: 292-293), thought is richer than language, and it is "impractical or even impossible to express exactly what one thinks".

Another problem that attends a postulation that thought is conducted in language is: Which language is it to be conducted in? A person who speaks more than one language can have a fully-formed thought without needing to decide which language to express it in. This means that their thought must be independent of the possible languages they might express it in. Jackendoff (1996, 1997, 2011) brings this point up as well: "If different languages can express the same thought, then thoughts cannot be embalmed in the form of any single language; they must be neutral as to what language they are expressed in…We would like to be able to say [bilinguals'] thoughts are essentially the same, no matter which language they are "thinking in." This is possible only if the form of thought is *neither* of these languages." (1996: 6/1997: 183).[13]

---

i) The assumption that the central epistemic notion is indeed the same across near-equivalent expressions at the level of mental representation. Even if treating them as the same poses no problems for linguistic description, it does not follow that their mental representations must contain the same shared component.

ii) The assumption that (even if these expressions are equal at the fine structure level in the epistemic qualification aspect), it is not possible to simultaneously conceptualize all the semantic aspects (epistemic modality, (inter)subjectivity, etc.) 'packaged' into the relevant expressions (If one were thinking in language, (assuming Nuyts' analysis) that is how the thought would be structured). See the discussion in Section 4.1 (Compartmentalization).

This is not to say I think Nuyts is wrong: I think he is right in the end—just that his argument here is insufficient. Nuyts' analysis is discussed in greater detail (and incorporated into our theory) in Section 4.1.

[13] An important caveat here is that for this argument to be successful, it must be realized that while in principle different languages can express a particular thought, they will do so with varying degrees of facility and faithfulness. This is apparently *not* the way Jackendoff envisions it (see the discussion in Section 2.1.4 on conceptual structure and linguistic structure below). It should be noted that without this caveat (i.e. with the assumption of absolute translation-equivalence), the argument can be used in the service of supporting the language-dependent position, as a defense against the assumption of implausibly strong Whorfian effects (determination of thought by the language one uses) that would otherwise attend it (c.f. Carruthers 2002: 61).

Another common argument (Pinker 1994; Levinson 1997, 2003; Vicente and Martínez-Manrique 2005, 2008; also c.f. Fodor 2001) is based on underdetermination of meaning in linguistic expressions. The idea is that expressions are ambiguous where the thought presumably is not; the mental representation of the thought, therefore, cannot be constituted by the ambiguous set of symbols that comprise the expression. This can be defended against by postulating that the representation of the expression that would constitute the thought would be *interpreted*, through contextual information, or that ambiguity in the expression reflects schematicity in thought (Frankish 1998; Gauker 2011, 2018). I discuss the issue of schematicity in language and thought in more detail in Section 2.2.1.3; I argue that schematicity is problematic in both directions: the expression can be more schematic than the thought, and the thought can be more schematic than the expression (both cases, of course, represent a mismatch between thought and language).

Further, the notion that because people generally are aware of their thoughts as linguistic expressions, it may be assumed that language is the medium of those thoughts, which underlies much of the thought-is-language position, has been criticized (Slezak 2002a, 2002b; Machery 2005) as being fallacious. Specifically, the expression of a thought is not necessarily the same thing as the thought itself, and it is a mistake to conflate the two. Jackendoff (1996, 1997, 2002, 2007, 2011, 2012) and Slezak (2002a, 2002b) have referred to this intuitive or introspective perception that thought is conducted in language as an illusion. In Jackendoff's view, thought itself is independent from language and unconscious; the phonological imagery of the linguistic expression of the thought, however, is available to consciousness. When a thought is accompanied by a phonologically-represented linguistic expression of it, the expression is confused for the thought: this creates the illusion that the thought is 'in' language. Jackendoff is careful to distinguish the thought from the expression: "phonetic form is not the form of thought; it is rather a consciously available expression of the thought" (1997: 187).[14]

[14] Note that the mechanics of this—the phonological representation providing an image that makes it available to awareness—are the same as that postulated by Carruthers (2002); but Jackendoff takes the opposite view regarding whether the expression is the same thing as the thought.

As far as introspection goes, at any rate, it does not always result in the impression that thought is conducted in language. Magee (1997) remarks on that proposition with incredulity: "I understand the view: what I do not understand, and have never been able to understand, is how anyone can hold it. For I find (and I do not believe that I am differently constructed from others in this respect) that it is directly contradicted by my immediate experience…always there is an act of "putting it into words" of which I am directly aware as being such…And it is always, of necessity, an inadequate rendering…how is it possible for so many undoubtedly clever people to assert, and believe in all sincerity, [that thought is conducted in language]?…How can they possibly believe what they say? What must their inner lives be like?" (pp. 78-86). The fact that people regularly have hard-to-express thoughts (mentioned at the beginning of this subsection) is an introspective basis for the opposite conclusion—that thought is independent of language. Hurlburt and colleagues (Hurlburt and Akhter 2008; Hurlburt et al. 2013), studying inner mental experience, find that subjects report experiencing what they call *unsymbolized thinking*, characterized as the experience of an explicit thought that "does not include the experience of words, images, or any other symbols" (Hurlburt and Akhter 2008: 1364 abstract)[15]. This experience is not universally reported (~25% of subjects have no unsymbolized thinking), and there is also individual variation in the prevalence of unsymbolized thinking (ranging from 0% to 80% of sampled experiences), which may reflect differences in intuition about whether thought is conducted in language or not. The experience of unsymbolized

[15] Hurlburt and Akhter provide examples of unsymbolized thinking like the following:
"Abigail is wondering whether Julio (her friend who will be giving her a ride that afternoon) will be driving his car or his pickup truck. This wondering is an explicit, unambiguous, ''thoughty" phenomenon: it is a thought, not a feeling or an intimation; it is about Julio, and not any other person; and it intends the distinction between Julio's car and truck, not his van or motorcycle, and not any other distinction. But there are no words that carry any of these features—no word ''Julio", no ''car", no ''truck", no ''driving". Further, there are no images (visual or otherwise) experienced along with this thought—no image of Julio, or of his car, or of his truck. In fact, there are no experienced symbols whatsoever—Abigail simply apprehends herself to be wondering this and can provide no further description of how this wondering takes place." (p. 1364)

thinking could be—the phenomenon is evidently hard to describe, and Hurlburt and Akhter (2008: 1372) are explicitly agnostic about this—conscious awareness of language-independent linguistic thought[16].

2.1.3 Indirect influence of language on thought

Now, all of this does not mean that there can be no influence from the target language on the thought. There are a number of ways that indirect influence can presumably occur.

---

[16] This is an appropriate place to remark on the conjectures of Jackendoff (1996, 1997, 2002, 2007, 2011, 2012) and Carruthers (2002) that phonological representation of (a linguistic formulation of) thought is necessary for conscious awareness of it. Their theses are:

> i) (unexpressed) thought itself is always unconscious (Jackendoff only; Carruthers makes no distinction between thoughts and their expressions)
>
> ii) conscious awareness of thoughts/expressions is possible only by "experiencing phonological images associated with them" (Jackendoff 2007: 83) (Both Jackendoff and Carruthers)

I reject both of these, though I do acknowledge that there are likely individual differences regarding what is conscious and what is not. On the first point, I think that binding a thought to linguistic structures is not required for awareness of it: see the discussion on conceptual lexicalization below (Section 2.2.3) and (probably) unsymbolized thinking above. Concerning the second point, as can be clearly seen in the 'tip-of-the-tongue' phenomenon, where a lexical item is accessed—and its (full) phonological form is not—(which Jackendoff himself brings up (2002: 217, 2007: 82)), the linguistic form can be directly accessible to conscious awareness without its phonological representation (see also Section 2.2.3). (Also, of course, if the first point is refuted, the second is automatically refuted as well.) Vicente and Jorba (2019) propose an interpretation of unsymbolized thinking that explains it as linguistically-expressed thought (i.e. the content of the "unsymbolized thought" is specific to a particular expression) with no phonology, which would be consistent with thesis i) but not with thesis ii). However, this does not appear to be coherent (if my understanding of the phenomenon of unsymbolized thinking is correct), since an expression—irrespective of any lack in phonological representation—necessarily *is* symbolized. See Section 2.3 (The architecture of language) for related discussion. All of this said, the phonological form may indeed facilitate awareness.

Language is a major vehicle for concept acquisition: people are exposed to concepts presented in the form of linguistic expressions—in some cases consisting of single lexemes and in others rather long-winded explanations—, and thus language may affect the concepts used in thought to a degree. This does not mean, however, that concepts as acquired and as involved in thought are strongly determined by language. First of all, this exposure is not deterministic of the concepts used in thought. Additionally, the concepts suggested by this sort of exposure are often not lexical concepts or combinations of them; the concepts and ideas referred to in actual discourse are often 'described' obliquely by the expression. Presumably, depending on how oblique in denoting or capturing the concept being gotten at the expression is, the strength of the tendency of this influence will vary. In the realm of communication (which is what is relevant to this particular point), expressions are formulated in such a way that the hearer or reader can interpret the intended message; this interpretation generally involves inference (Levelt 1989; Sperber and Wilson 1995; Wilson and Sperber 2002; Carston 2002; etc.)[17], which can arrive at an understanding that is rather different in form from what was explicitly said. The more oblique the denotation, the greater this difference tends to be—and the less determining the expression itself is of the acquired concept. And, of course, conversely, the more 'direct' the denotation, the greater the degree to which the concept tends to be shaped is.

Similarly, linguistic 'handles' can be used to access concepts at the level of thought, either in a discourse or simply as a 'set concept'. In a discourse, ideas that the various participants express are referred to by others. In order to formulate a thought containing a reference to another's idea, that idea must be mentally represented; and since the exposure to that concept was through a linguistic expression of it, this mental representation can be influenced by language. This is effectively similar to the concept acquisition case just mentioned above. This mechanic applies as well to one's own expressed ideas: their original formulation may be language-independent, and, after they are expressed in language, later reference to them may recall them in a more language-dependent representation (e.g. if some idea is referred to twice—once before

[17] This inference includes abduction; however, abduction for interpretation is a different topic from that of this dissertation, so I will not discuss it.

linguistic formulation and once after—the second access may retrieve a linguistic representation that is an approximate equivalent of the first language-independent representation rather than that original representation itself). These effects do not require discourse per se: inner speech can give the same sort of effects. One case of this function of linguistic 'handles' is the use of language to "offload" ideas to reduce cognitive load, as has been suggested by, e.g., Jackendoff (1997: 200-202) and Clark (1998).

Another source is that seen in linguistic priming effects. For example, numerous experiments have shown that processing of a word can be made faster if subjects were exposed to a semantically- or phonologically-related word beforehand; also, exposure to a particular construction in one language can make production of that construction in another language more likely. While these effects may be limited to the accessing of linguistic structures, another possibility is that they involve accessibility of certain structures in thought. Linguistic priming effects will be discussed later in Section 4.2.2 (Accessing elements).

Additionally, there may be some tendency for people to think things that can be expressed. Some ideas are harder to express in (a particular) language than others, and this may prompt 'intentionally'—but not necessarily consciously—thinking in an at least somewhat compatible form, at least sometimes (c.f. Levelt 1989; Slobin 1996).

All of these effects may result in habitual patterns of thinking aligning with language more than they would by pure chance (c.f. Slobin 1996; Levinson 1997). So the original thought may tend toward some degree of similarity with potential linguistic expressions.

2.1.4 Linguistic structure and conceptual structure

So, we have established that thought is generally independent of language, but can receive influence from it. The next topic we examine is the relationship between linguistic structure and conceptual structure. If thoughts are non-identical with their linguistic expressions, there are two versions of conceptual structure we must consider: the conceptual structure of the original thought and the conceptual structure expressed

in language. Differences between what is thought originally and what is expressed are fundamental to the subject we are looking at; this is the problem we take up now.

2.1.4.1 Views on conceptual structure

We will discuss two opposing 'sides'/facets of the relationship between language and conceptual structure: conceptual structure is in thought, and conceptual structure is in language.

Very broadly speaking, researchers in generative syntax consider conceptual structure to be separate from linguistic structure, and those in cognitive linguistics consider conceptual structure to be identical with linguistic meaning. The former group tends to equate linguistic structure with syntax, and the latter group tends to equate conceptual structure with semantics. The divergence in these views seems to come more from a difference in perspective resulting from what is to be emphasized in the respective objects of study than a substantive disagreement. My intention here, however, is to discuss these views for their own sake as complementary facets, rather than as opposing factions' views.

Firstly, conceptual structure, as a matter of obviousness, must be part of thought. Language is in general insufficient in capturing this conceptual structure (see Jackendoff 1996: 28-30/1997: 206-8; Levinson 2003: 295; etc.). So on one hand, conceptual structure is a matter of thought, and separate from language.

The current mainstream generative position of Chomsky and colleagues concerning conceptual structure, meaning, and language, advanced in Hauser et al. (2002) and others[18], divides the abilities involved in language use into three systems: the sensory-motor system, the conceptual-intentional system (also referred to in Berwick and Chomsky 2011 as the "system of thought" (p. 27) and the "semantic" system (pp. 33-34)), and the computational system for syntactic structure. Having assumed this architecture, they hypothesize that only the ability of syntactic

[18] These researchers are concerned primarily with issues of architecture and the language faculty; their ideas will be discussed in greater detail in Section 2.3 (The architecture of language).

computation is uniquely human (although other animals have conceptual systems, "the conceptual resources of humans that enter into language use appear to be far richer" (Berwick and Chomsky 2011: 39); also see Berwick 2011). They suggest that evolutionarily, the conceptual system and the sensory-motor system necessary for externalization existed first, and it is the emergence of the ability to compute syntactic structure that linked them and enabled language (particularly pp. 1572-1573, 1578). The interpretation of linguistic expressions is effected by a mapping from syntactic representations to semantic representations. In order to do this, syntactic representations are interpreted by the conceptual system (or at the interface with the conceptual system—Chomsky and colleagues are vague and possibly somewhat inconsistent on this point[19]). Thus, conceptual structure is distinct from linguistic representations.

It should be noted that a common (sup)position in the generative tradition (e.g. Bickerton 1995; Hinzen 2006, 2011; Chomsky 2007; Berwick and Chomsky 2011) is that the mechanisms that enable the use of language also enable thought[20]: an abstract language of thought can be implemented by the same cognitive operations that are responsible for implementing language. This is predicated on the idea that thought takes place in an abstract language that has essentially the same fundamental syntactic structure of actual language. This is not a proposal that the thought corresponding to

[19] Hauser et al. refer to a syntactic computational system that "generates internal representations and maps them into the sensory-motor interface by the phonological system, and into the conceptual-intentional interface by the (formal) semantic system" (2002: 1571); Berwick and Chomsky talk of syntactic expressions being "interpreted by conceptual systems" to yield a language of thought (2011: 30) and of "a computational system efficiently generating expressions interpretable at the semantic–pragmatic interface" (2011: 37).

[20] And thus, the evolutionary emergence of language would have resulted in a much more sophisticated, uniquely human form of thought.

particular expression is the same as that expression[21]; rather, it is a proposal that the two representational formats (thought and language) have fundamentally the same structure. The motivation for this idea is that the combinatorial and hierarchical syntactic structure of language presumably is shared by an abstract language of thought. In the Minimalist Program of generative syntax research undertaken by Chomsky and colleagues (following Chomsky 1995), syntactic structure is (somewhat tentatively) held to be generated by essentially one operation, Merge[22]; the cognitive ability to merge two cognitive objects is responsible for structuring both thought and mental representations of linguistic expressions.

Jackendoff (1996, 1997, 2002, 2007, 2011, 2012) advances a somewhat different conception. Although Jackendoff "take[s] issue with every part of [Hauser et al.'s] conception of the language faculty" (2011: 591), his theory is similar in that it posits separate components—phonology, syntax, and semantics—that interface with each other. (See Section 2.3 (The architecture of language) for more discussion.) In his view, conceptual structure "is not part of language per se—it is part of thought" (2002: 123). It is independent from syntax and phonology, and contains the representations and combinatoriality with which thought is conducted; it is "largely autonomous of language" (in keeping with Jackendoff's view—discussed earlier—that thought is independent of language) (2007: 192-3). At the same time, conceptual structure is held to be identical with the semantic structure of the linguistic expression. The semantic

[21] Hinzen may be an exception to this (it seems he thinks thought is conducted in language), though he is not explicit on the point, beyond remarks such as "If I am right, however, that syntax gives us [the conceptual-intentional system] as we know it, rather than answering any conditions it imposes, there is, strictly speaking, no such thing as a mapping to the semantic interface." (2011: 426). He states, concerning the different "referential perspectives" possible for a particular piece of content, (e.g. *Caesar's destruction of Syracuse* / *Caesar destroyed Syracuse* / *(That) Caesar destroyed Syracuse*), that "grammar, through the functional layers of lexical projections, provides the perspectives in question, which nothing in the external world or in non-linguistic cognition seems to determine" (2011: 434). So deciding on an expression determines the conceptual structure. The emphasis on conceptual structure as expressed in language is shared by the cognitive linguistics perspective discussed below.

[22] Chomsky distinguishes between *external merge* (responsible for embedding) and *internal merge* (responsible for displacement).

structure does not depend on language: all that changes depending on language is the syntax and phonology: for example, "[w]hat it means to translate from English to French is to take the semantic/conceptual structure of an English sentence and clothe it in the syntax and phonology of French (including French vocabulary)" (2007: 83). That is to say, simply, that he takes the meaning of the expression to be what was intended.

On the other hand, linguistic expressions depict conceptual structure. (Cognitive linguistics generally makes it its aim to analyze this conceptual structure.) To give a classical example, 'the glass is half full' and 'the glass is half empty' are/depict different conceptual structures resultant from different construals of the same situation. This other position emphasizes conceptual structure as expressed in language[23]. Research in this paradigm (e.g. Lakoff and Johnson 1980; Langacker 1987; Talmy 2000a, 2000b) generally aims to understand aspects of cognition seen through the linguistic expression of conceptual structure.

One such aspect is *mental scanning* (Langacker 1997, 2008, etc.). Mental scanning in construal can be seen in, for instance, *There is a house every now and then through the valley* (as contrasted with *There are some houses in the valley*); the former represents a construal where the group of houses is 'scanned' as an array and apprehended as a figurative sequence, while the latter does not involve any scanning (Talmy 2000a: 71)[24]. Mental scanning can involve specific directions: for example, *The roof slopes steeply downward* and *The roof slopes steeply upward* show different directions of mental scanning (Langacker 1997: 244).

Another is *reference point relationships* (Langacker 1987, 1993, 2008, etc.). For example, the two expressions *Your camera is upstairs, in the bedroom, in the closet, on the shelf* and *Your camera is on the shelf, in the closet, in the bedroom, upstairs* (2008: 81) refer to the same location, the first by "zooming in" (*upstairs > in the bedroom*…), and the second by "zooming out" (*on the shelf > in the closet*…). In doing so, locations

[23] This, of course, does not mean that the existence of non-linguistic conceptual structure is denied. See, for example, Langacker (1987: 97-98) and also the discussion in Section 2.1.4.2 (Non-linguistic structure) below.

[24] Talmy refers to this as *sequentializing*. I am using the more general term *mental scanning* (c.f. Langacker 1997: 244, also citing Talmy's example as showing mental scanning).

are identified in reference to other locations (i.e. *–Where upstairs? –In the bedroom.*, etc. or *–Which shelf? –The one in the closet.*, etc.). The two expressions differ in which locations are used to refer to which. The use of reference points in conceptualization is quite broad; another simple example can be seen in *Do you see that boat out there in the lake? There's a duck swimming right next to it* (2008: 83).[25]

An additional major sort of conceptual structure that can be seen in linguistic expressions is *metaphor*. In conceptual metaphor theory (Lakoff and Johnson 1980, 1999; Lakoff 1987, 2008; Johnson 1987, 2007; Sweetser 1990; Grady et al. 1996; etc.), metaphorical expressions are taken to reflect the conceptualization of one thing in terms of another. A large portion of linguistic expressions can be identified as metaphorical; typical examples of metaphor include MORE IS UP (e.g. *a higher amount*), AFFECTION IS WARMTH (e.g. *a warm welcome*), and temporal metaphors (see especially Evans 2003, etc. and Moore 2014, etc.) like *Christmas is coming*. The claim of conceptual metaphor theory is that in, for instance, *a warm welcome*, *warm* refers firstly to physical warmth, and this, through the associative mapping of AFFECTION IS WARMTH, yields the metaphorical meaning (an "affectionate" welcome).

These are only a few aspects of the conceptual structure depicted in linguistic expressions. These and other aspects will be discussed further in Section 2.2.1.4. Essentially, the point of language is to depict conceptual structure.

Both of these prototypical stances are valid in their own right: certainly, conceptual structure can be seen in language; and certainly, thought is independent of language and linguistic expressions are not perfect descriptions of thought, and thus conceptual structure is separate from language. Neither, however, is suitable for constructing an appraisal of the problem of expressing thought in language; rather, we must adopt both of these contradictory stances as once. That is: there are two different versions of conceptual structure that must be considered; one is the conceptual structure as originally conceived in language-independent thought, and the other is the conceptual structure as expressed in language.

[25] This sort of conceptualization can also be applied to metaphorical uses. See Section 2.2.1.4 below.

The basic problem with the sort of segregated understanding proposed by Hauser et al. (2002) and others is that it does not facilitate a mechanistic understanding of the connection of thought to linguistic meaning. A more substantive link between thought and language is needed. While it is important to note that if one were to adopt this stance, it would still be possible to formulate the conclusion of the theory described in this dissertation in its most basic form—that expressing thought in language is abduction—it is not possible to describe the problem in detail using this approach.

In Jackendoff's view, linguistic structure can change irrespective of conceptual structure, so conceptual structure is separate from linguistic structure, even as it is part of linguistic structure (i.e. the semantics). It is important to note that what Jackendoff means when he equates meaning with conceptualization is that he takes the philosophical position that meanings refer not to objective things that exist in the 'real world', but to things in "the world as human beings conceptualize it" (2007: 192), so this is a much less self-contradictory stance than it may initially seem. But it is still not tenable. Having semantic structure be determined by conceptualization and conceptual structure (which is equal to the semantic structure) independent of the target language is not coherent. Different languages express the "same thing" in different ways (see Chapters 3 and 4 for a detailed exposition of this phenomenon), so the semantic structure inherent in the various expressions of the same thought is necessarily different: different expressions of the "same thing" express different conceptualizations. This means that having conceptual structure equal to semantic structure requires that it be dependent on the particular expression[26]. But having conceptual structure be independent of linguistic structure requires that conceptual structure be independent of the particular expression. Now, both statements are true: the meanings of expressions reflect conceptualization, and conceptual structure (in general) is language-independent. The problem comes from conflating two different versions of conceptual structure (that of the original thought and that as expressed in language). This imprecision is the

[26] This, of course, would not be an issue if one posits that thought is done in the target language—but Jackendoff explicitly argues against this (as do I).

source of the odd notion that semantic structure does not vary between languages.[27] So, while, as Jackendoff remarks, French, Turkish, and English speakers can have the same thoughts, when they are expressed in language (i.e. put into French or Turkish or English), the meanings of those expressions are necessarily different (though still, obviously, mostly similar). As Jackendoff himself notes, the expression of the thought is not the same as the thought itself. Recognition that the semantic structure of expressions does change between languages (and, as a matter of fact, within individual languages as well) is crucial for our purposes.

The 'alternative' view discussed above—that of cognitive linguistics—certainly takes differences in semantic structure between expressions as a central object. But this view also has a fundamental problem. With this perspective, one can only consider the conceptual structure as represented in the expression: but the fact that an expression represents a particular conceptual structure does not mean that the original thought must have that same conceptual structure. Now, this is still true in some sense: no one would ever argue the point that someone would say (in earnest, at least) *the glass is half empty* and not *the glass is half full* unless they think that it is half empty and not half full. Likewise, to reference an example of Levinson's (1997, 2003), surely no one would say that some object is *north* of some other object and not that it is *to the left* of it without actually thinking that it is north of it. As Levinson (1997: 15) remarks, as a trivial point, anything we say we think. But if thought is independent from language, then the

[27] It should be noted that this also affects his argument from translation-equivalence for the independence of thought from language. His argument (1996: 6, 1997: 183, 2007: 83, 2011: 613), in essence, is that different languages can express the same thought, and therefore thought must be independent of language; this is perfectly sensible, and is, in fact, one of the arguments that I myself make (see Section 2.1.2 above; but c.f. footnote 13). However, conflation of the meaning of the expression with the originally-intended meaning turns the reasoning into:

-The intended meaning can be the same. (True)

-The expressed meaning is the same, regardless of language. (False)

-Therefore: i) Thought is independent of language. (True, but now fallacious)

ii) Linguistic structure is independent of meaning. (False)

See Barnden (1996) for a similar critique (but one framed—at least for the sake of argument—as *This does not rule out thought being conducted in language*) of this point.

conceptual structure represented in the expression is not, in general, going to be the same as the conceptual structure of the original thought (at least at a sufficiently fine level of detail). Different languages and different situations or linguistic contexts within languages differ in what can be expressed, so even if—by luck—there exists some specific case where the conceptual structure depicted in a particular expression is a perfect match for the conceptual structure of the original thought, the general case is that the two conceptual structures will diverge to some degree.

So, conceptual structure features both in language-independent thought and in mental representations of linguistic expressions. In general, these two versions of conceptual structure will be different, but quite close. Linguistic expressions *approximate* thought. This theme is of importance throughout the dissertation.

2.1.4.2 Non-linguistic conceptual structure

In the above subsection, we have been discussing *linguistic* conceptual structure. Now, as not infrequently noted (Pinker 1994; Gauker 1994; Keller and Keller 1996; Jackendoff 1996, 1997, 2007; Barnden 1996; Levinson 1997; Tomlin 1997, etc.), there are other sorts of mental representations that are not properly 'linguistic'. Perceptual imagery (perhaps most obviously visual imagery, but also including auditory imagery, tactile imagery, proprioceptive imagery, etc.) and various kinds of artistic thought (e.g. for musical composition, painting, sculpture, etc.) all involve kinds of non-linguistic conceptual structure. These kinds of representations are of a different sort from the mental representations of linguistic expressions and of language-independent linguistic thought (i.e. thought that can be expressed in language). Despite the conceptual structure involved being non-linguistic in this sense, it is not entirely unconnected to linguistic conceptual structure and linguistic expressions. In this subsection, we will consider the relation of non-linguistic conceptual structure to language. We will discuss two aspects/means of connection: conversion and evocation.

The first concerns the rather obvious point that non-linguistic conceptual structure can be converted to linguistic conceptual structure. One can, for example, behold a painting and obtain an image of it. Now, this image is not necessarily purely visual—it

can be an affective or artistic representation of the painting. At any rate, this representation is non-linguistic. But one can describe this representation to someone else, essentially 'translating' it into language. One can say there is a river in the middle, with a tree on the other side and a cute furry animal in the foreground; one can say that the cute furry animal is a symbol of heroism or nihilism or whatever; one can say that part of the painting conveys unbridled joy and another part lethargy; and one can say that the painting as a whole captures the essence of the human condition. That is: non-linguistic conceptual structure can be converted to or used to derive linguistic conceptual structure. It is not clear how this case relates to our problem. Very simply, it is possible that explanation of non-linguistic conceptual structure in language involves essentially the same process as expressing linguistic thought does (in which case our theory would apply to it as well), and it is also possible that it instead is fundamentally different, and what is expressed in language is not the non-linguistic conceptual structure itself, but a different linguistic thought derived from it (in which case our theory would apply only to the expression of this linguistic thought). At any rate, this issue will not be a focus of our investigation.

The second depends on the idea that non-linguistic conceptual structure can be evoked by linguistic expressions. One understanding of the function of language in communication is that expressions serve to evoke mental representations, rather than representing meaning directly (e.g. Sperber and Wilson 1995; Carston 2002). In simulation semantics accounts, this evoked mental representation is postulated to be of a non-linguistic sort. Barsalou et al. (2008) and Evans (2009, 2016) propose that language does not directly encode conceptual structure, but rather encodes schematic representations that evoke imageable simulations (in the sense of Barsalou 1999). The content of this simulated experience is perceptual and non-linguistic[28]. For example, the expression *red ball in a box* would evoke a visualization of a red ball in a box (Evans 2009: 108). Since linguistic representations cannot directly encode the perceptual experience itself (see Evans 2009: 105), Evans distinguishes between what he calls *linguistic content* (i.e. what is actually represented by linguistic expressions) and

[28] See Section 2.2.2 for a discussion of perceptual representations and their relationship to conceptual representations and Section 3.3.3 for a discussion of simulation.

conceptual content (i.e. the perceptual representations that form the basis for the simulations that can be evoked by linguistic content). Conceptual content, in Evans' account (based on the conception of concepts advanced by Barsalou (1999, 2003, 2017, etc.)), is comprised of associatively structured collections of perceptual images—i.e. perceptual encyclopedic knowledge[29]; the role of linguistic content is to "afford access" to this conceptual content. This accounts for some of the flexibility seen in language use. For example, the word *red* evokes different perceptual images—i.e. different kinds of red—in *The teacher scrawled in red ink all over the pupil's homework exercise* and *The red squirrel is in danger of becoming extinct in the British Isles* (ibid.: 206). So in this view, the relationship is that linguistic structure *evokes* conceptual structure, but does not *encode* it: i.e. linguistic expressions are not held to represent conceptual structure (in contrast to the view of Langacker, Talmy, and others discussed above in Section 2.1.4.1), but rather to suggest it indirectly (Evans 2009, particularly Ch. 6, 186, 189-192). It must be noted that this is a theory of comprehension (see Evans 2009: xii), not the production we are concerned with. Certainly, simulated experience can be a form of mental representation; when it constitutes content that is to be expressed in language, it would be an instance of the case of conversion from non-linguistic conceptual structure discussed above.

2.1.5 Synopsis

We can now explicitly indicate the stance taken in this dissertation.

Language expresses linguistic thought. Linguistic thought can be expressed in language, but it does not have to be. Conceiving a linguistic thought does not entail 'translating' it into language. But every linguistic expression is based on some linguistic thought.

[29] Encyclopedic knowledge is open-ended functional and associative information about conceptual objects. See, e.g., Haiman (1980) and Langacker (1987, 2008) on the role of encyclopedic knowledge in linguistic meaning. This will be discussed more in Chapter 3 (Meaning Construction).

Thought is, in general, language-independent. It is not restricted to the specific meanings available in particular languages, and a particular thought can be expressed in any language (albeit, perhaps with varying degrees of difficulty and success). However, each expression will have its own way of expressing it: each will have a unique way of constructing its meaning. So while translation-equivalents between languages and paraphrases within languages may be suitable expressions of a particular thought, what they actually express will differ. The meanings that are actually able to be expressed depend on language and on linguistic context within a particular language. The fine-structure level of meaning of a particular expression will in general be a property of that expression, and not of the underlying thought itself. Thought is not, in general, conducted using the concepts and meanings available in a particular language.

That said, there are caveats to this notion of language-independence that should be noted: concept acquisition is often done through the medium of language, which may influence the concepts acquired; one can reference (in thought) ideas explicitly expressed in language; linguistic structures may produce priming effects in thought; thought can be adapted to the requirements of the target language (thinking for speaking); and these indirect influences may affect habitual patterns of thinking even when expressing thought in language is not an immediate concern. So 'language-independent' thought probably has a greater-than-chance resemblance to potential linguistic expressions.

Linguistic expressions depict conceptual structure, but this conceptual structure is, in general, not identical to that of the language-independent thought. In principle, the linguistic expression is the best possible approximation to the original thought. However, differences in mental accessibility present a perturbation to this ideal: higher-frequency linguistic elements tend to be more easily accessed (Alario et al. 2002; Arnon and Snider 2010; Janssen and Barber 2012; etc.), and various priming effects can make certain linguistic structures more accessible as well (Meyer and Schvaneveldt 1971; Bock and Loebell 1990; Loebell and Bock 2003; Hartsuiker et al. 2004; Bock and

Ferreira 2014; Goldrick 2014; etc.); more mentally-accessible expressions may be chosen over others that better approximate the thought in a strict sense[30].

So, the start-point of our investigation has been established.

2.2 Linguistic Meaning and its Mental Representation

Next, we will consider two separate but related issues. One is the question of what constitutes the meaning of linguistic expressions, and the other is the question of what the format of their mental representation is. I do not think it is possible at this point to deliver a definitive or conclusive answer to either, and I will not endeavor to do so. However, I do believe that the questions can be profitably explored. I hope to provide, at least, what may amount to the beginning of an understanding of what the issues are in the discussion below.

2.2.1 The nature of linguistic representations

2.2.1.1 The (non-)distinction between lexicon and grammar

Our first topic is how to treat grammar and the lexicon. Linguists are divided on this issue: some treat the lexicon and grammar distinctly, and others do not.

There are two principle bases for a dichotomous treatment of grammar and the lexicon.

One is the view that syntax is (or at least can be treated as) autonomous from meaning. If, as in generative syntax, the object of study is to come up with a set of rules that formally describe grammatical well-formedness, it makes sense to distinguish between grammaticality and semantic coherence: an expression can be semantically interpretable but ungrammatical (see e.g. Chomsky 1957; Newmeyer 1998).

[30] Thanks to Zheng Xinshuang (personal communication) for making me aware of this. Priming effects will be discussed again in Section 4.2.

The other is the notion that there is a functional division between the lexicon and grammar. This is an important component of the theory advanced by Talmy (2000a, 2000b). The idea is that there is fundamental division between the structure and content of the meaning of linguistic expressions. Lexical items provide the content, and grammatical items provide the structure. Lexical items are characterized by being what he calls *open-class*, which means they are a not restricted set that can be added to more or less freely. By contrast, grammatical items are characterized by being what he calls *closed-class*, which means they are part of a restricted set that cannot be added to easily. It is possible to come up with a new lexical word more or less at will, but cases, determiners, etc. are not so accommodating. Grammatical meanings are also restricted in the sorts of meanings they can express. For example, whether something is singular or plural can be expressed grammatically, but whether something is red or blue cannot be, apparently in any language (Talmy 2000a: 24). One merit of this sort of treatment is that it cleanly accounts for the fact that only some kinds of meanings are grammaticalized.

However, there are several related reasons that a dichotomous treatment of the lexicon and grammar is problematic.

One is that grammatical elements also contribute meaning—i.e. their role is not restricted to content-less structuring (c.f. Langacker 1987, 2000, 2008, 2009b etc.). For example, number (singular/plural/dual etc.), tense, definiteness (e.g. *a* vs. *the*), etc. are fundamentally aspects of meaning. It is important to note that this includes not just "function words" and grammatical morphology but constructions as well—they are a vital and ubiquitous source of meaning (c.f. Goldberg 1995, 2006; Croft 2001; Langacker 2009a; etc.). A commonly-given example is the "caused-motion" construction in English (*She tossed the book to him*; Goldberg 2009: 97 / *He sneezed the napkin off the table*; Goldberg 1995), where imparting of movement is expressed by the construction. Similarly, the construction in *that travesty of a theory* gives it a meaning something like 'that theory, which is a travesty' (Jackendoff 2011: 610). Croft (2012: 15) notes that *But reading to a dog isn't so scary* and *When you're 6 or 7 years old, that's quite a lot of dog bearing down on you* can be described as involving a "count noun" construction and a "mass noun" construction, respectively; these give different

meanings. In certain situations, Hebrew expresses 'be' with a copula, and in others expresses it with pure syntactic arrangement, as in *achshav ani ha-sho'er* (lit. 'now I the-goalkeeper; 'Now I am the goalkeeper') or *zot ha-siba* (lit. 'that the-reason'; 'That is the reason') (Glinert 1989: 171); c.f. *hayiti ha-sho'er* (lit. 'I-was the-goalkeeper') and *hayta ha-siba* (lit. 'that-was the reason').

Another is that, as Croft (2007: 344) points out, grammatical meanings and lexical meanings are not cleanly separable; many sorts of meanings expressed by grammatical elements can also be expressed by lexical elements. For example, pastness can be expressed grammatically, by a tense inflection (*When he arrived*) or lexically (*On his previous arrival*) (Talmy 2000a: 35/Croft 2007: 344). It is often possible to express the sort of information that grammar can encode with a lexical item.

Another issue is the phenomenon of grammaticalization[31]; there are two aspects of the problem that this presents (Croft 2007: 344-346). One is the fact that there can be transitions from lexicon to grammar (and these happen continuously, not saltationally), which means that the sorts of meanings that lexicon and grammar have are compatible—if they had different sorts of meanings, a continuous transition would not be possible. The other aspect of this is that at any particular time, there are structures at various points of transition. There are some that are fully lexical, and others that are fully grammatical, and others that are at various intermediate stages of grammaticalization.

For these reasons, many researchers posit a lexicon-grammar continuum (e.g. Langacker 1987, 2008; Goldberg 1995; Croft 2001, 2007; Jackendoff 2002, 2011; Evans 2009). This does not mean no distinction can be made, of course. In Croft's view, while a dichotomy between grammar and the lexicon is to be rejected, the content-versus-structuring paradigm is still appropriate in some sense: structuring in expression

[31] Grammaticalization is the process of establishing a linguistic structure (which is originally not grammatical) as a grammatical element. A standard example is *be going to* as a (grammatical) expression of futurity being derived from a (non-grammatical) expression of physical movement (and not of futurity). Another is *have* as an expression of possession being grammaticalized to being used to indicate perfectivity, as in *She has a cat* vs. *She has finished* (Langacker 2008: 87). See, for example, Heine et al. (1991) and Hopper and Traugott (2003) for detailed examinations of this phenomenon.

does indeed tend to be done through grammar; the means by which structuring can be done are relatively limited, and these tend (by the resultant high-frequency use) to become grammaticalized. Langacker (e.g. 2008: 21-23) and Evans (2009) hold that grammatical meanings are more schematic, and Evans says a distinction can be drawn on the basis of whether the element "affords access" to the conceptual content from which a perceptual simulation can be built.

For our purposes—describing the expression of thought in language—, grammar is not a set of rules to be followed; it is part of the means to make meaning (c.f. Jackendoff 2002: 289-291; Langacker 2008: 23-24; Evans 2009: 36). Grammatical rules, 'real' as they are as accurate descriptors of what is well-formed and what is not, are simply reflexes of (some of) the mechanisms used to express meaning in language. Since lexical elements and grammatical elements are clearly not dichotomous in function, it does not matter if a particular linguistic element is "lexical" or "grammatical" in that, fundamentally, it is just a potential tool that can be used to express meaning. This is a main reason that I use the term *linguistic element*: it gives the necessary generality in covering both encoding of "syntactic" information and "lexical" information with no principled difference. This does not mean that there can be no identifying of a particular structure as grammatical or syntactic as opposed to lexical, but it does mean that lexical aspects and grammatical aspects are not to be treated in isolation, as they work together to build the meaning of the expression (c.f. Bock and Ferreira 2014). So I will treat lexical aspects and grammatical aspects of linguistic structure in an integrated manner.

2.2.1.2 What counts as an element?

As mentioned above, what I refer to as *linguistic elements* includes words, morphemes, and constructions. In the narrowest sense of the category, lexical items are words. However, it is often given a more generalized characterization; at the other extreme, it includes productive morphemes, phrasal verbs etc., idioms, fixed expressions, and constructions. This generalized conception can be defined alternatively as referring to the set of items "entrenched" in the speaker's mind (in Langacker's (e.g.

2008: 16-17) parlance) or the set of items stored in long-term memory rather than being constructed online in working memory (in Jackendoff's (e.g. 2002: 152-154) terms). So whereas *moon*, *moonless*, and *moonless night* generally qualify as lexical items in this sense, *dollarless* generally does not (Langacker 2008: 16-17). The notion of constructions is similarly often maximally generalized, and defined as any form–meaning connection (e.g. Croft 2001), which is essentially what I mean by linguistic element. Now, in various places throughout the dissertation, it will be valuable to distinguish between meaning encoded in words, morphemes, and constructions. A construction, as I refer to it, is a syntactic pattern that carries a particular meaning; the pattern may or may not be associated with particular words or morphemes (for example, [to X] (e.g. *It's not a big deal to me*) can be considered a construction (associated with *to*), and so can the transitive-verb–direct-object relationship expressed in the syntactic pattern [VP NP] (e.g. *eat an apple*) (not associated with any words or morphemes))[32]. I will refer to *words* or *morphemes* or *constructions* when specification from among them is relevant and *linguistic elements* when it is immaterial.

The main point I want to make in this section is that what counts as an element and what meaning each element is held to have are not fixed, but rather depend on the individual language user's analysis (c.f. Dąbrowska 2012). (A similar point is made in Section 2.2.1.3 (Granularity of linguistic meaning) below.)

Different analyses are possible for particular linguistic structures. The actual concepts associated with linguistic structures will differ (relatively minutely) from person to person; for example, different people will have different conceptions of 'beauty'. Structural aspects of the representation of meaning in linguistic expressions are also not fixed. Different analyses of Talmy's (2000b: 49) example *The bottle floated into the cave* concerning where movement is encoded have been proposed: movement can be seen as being carried by the verb (*float*) (Talmy 2000b), the construction (<VERB into X>) (Goldberg 1995), or both the verb and the construction (Langacker 2009a)[33].

[32] What I mean by 'syntactic pattern' is not just positioning of words. Constructions can also be determined by patterns of indexation (case marking etc.). This is particularly salient in languages with free word order, like Latin. See Section 4.1.2.1.3 for more discussion.

[33] This example will be discussed in more detail in Section 4.1.

Different language users can in principle assign movement to any of these encodings in their mental representation of the expression.

Different linguistic structures differ in how readily they are understood as being comprised of subcomponents (i.e. their *analyzability* [34]). For instance, at the morphological level, in some cases, a word is not seen as being composed of several morphemes; in other cases, different component morphemes can be identified, but their meanings are not clear or known; and in other cases, they are clearly known, and the meaning of the word can be understood as being built up from them. For example, the word *metaphor* is probably for most speakers not analyzable; other speakers may identify the two morphemes <meta> and <phor> (perhaps by association with other words, like *metaphysics* and *theophoric*), but not know their meaning, and others may know what they mean, and some may even understand the meaning of the concept as being represented by 'carry across'. Similarly, as Levelt et al. note (1999: 12), *replicate* formally/historically consists of two morphemes (coming from Latin *re* + *plicare* ('fold again')), but is generally analyzed as a single morpheme. Unless the language user is of a very Latinate persuasion (and might be inclined, for instance, to use *express* to mean 'press out'), the two-morpheme analysis is impossible.

As Langacker (2000: 152, 2017: 43; etc.) points out, there is a general gradience in analyzability, with cases like *squealer* more readily analyzable as composites than cases like *computer*, which in turn are more analyzable than cases like *drawer*, which have almost no analyzability. Ignoring or perceiving internal structure does not have to be all-or-nothing. For example, the Japanese *shinakerebanaranai* (lit. 'if not do, not become'; 'need to do') can be analyzed as an indivisible unit (particularly in its contracted form, *shinakya*, which cannot be inflected (c.f. inflection of the un-contracted form *shinakerebanaranakatta* ('needed to do'))), as two parts

[34] It is unfortunate that the term for a language user's general understanding of a linguistic structure and the term for specifically the decomposability of a linguistic structure overlap. In my writing here, generally only the form *analyzability* refers to the latter; other forms (e.g. *analyze*, *analysis*) refer to the former.

(*shinakereba*/*naranai* ('if not do/not become')), or be further decomposed (e.g. *shi*(/)*nakere*(/)*ba*/*nara*(/)*nai* ('do/not/if/become/not'))[35].

Entrenchment (Langacker 2002, 2008, 2009, 2017; Schmid 2015; etc.) can encourage analyses of composite structures as single elements. There may be some tendency for entrenched common phrases to be accessed as a whole without being built up from their constituent parts—i.e. stored and accessed as a whole, and thus constituting a lexical item in Jackendoff's sense (2002, above).

Now, it is probable that in some cases, routinely-used composite linguistic structures are accessed as complex units, where the constituent meanings are mentally represented as well[36]. So, for example, *in order to* may be accessed as a pattern consisting in multiple elements. There are a few possibilities regarding what should be considered an element in these cases. One is that the complex as a whole is the element; another is that the constituents are the elements. The distinction is whether (for the first) the meaning of the complex as a whole is stored as an object in its own right (i.e. it is an *indexed concept*; see Section 2.2.3), or (for the second) it is not, and while access of the pattern does exist as a routine, the complex as a whole is not directly associated with a particular meaning. It is also conceivable that both the unique meaning of the whole and the meanings of the constituents can be simultaneously represented in overlapping structures; in this case, both the complex and the constituents are elements. At any rate, I do not think that anything I have to say will be confusing on this matter.

The particular analysis, in addition to depending on the individual, can even vary depending on context, so a particular linguistic structure may be analyzed differently in different situations[37]. For example, *socioeconomic* may in some contexts be understood

[35] This example and analysis was provided by Akira Machida.

[36] Langacker refers to this as the *unit status* of a complex structure (e.g. 2008: 16-17). He notes: "It is important to realize that unit status does not entail the absence or unimportance of components, merely the routinized nature of their execution (which does not however tend to diminish their individual salience)."

[37] This is also an important part of historical change: If a novel situation-specific meaning is entrenched (i.e. established within an individual) and conventionalized (i.e. established within a language community), it will be directly associated with the linguistic element in the analyses of language users, either replacing the original meaning or existing alongside it.

as a conceptual composite (i.e. 'social and economic') and in others as a single atomic concept. In the following subsection, we will see more aspects that depend on the language user's analysis, which varies according to the individual and the situation.

2.2.1.3 Granularity of linguistic meaning

In this section, first we will address the question of how detailed linguistic meaning is, and then we will consider how the granularity of linguistic meaning compares to that of thought. Our object here, granularity, concerns what level of detail is represented in linguistic meanings. Differences in granularity are, alternatively, differences in relative vagueness.

The points that I wish to establish here are:

i) There is great variation in the granularity of linguistic meaning, and the specificity of linguistic meaning is not a 'stable' or fixed thing (as the line between polysemy and schematicity/vagueness is quite blurry).
ii) As has been remarked by many researchers (e.g. Pinker 1994; Jackendoff 1996, 1997; Levinson 1997; Evans 2009), linguistic meanings are often less precise than what may have been intended in the thought; but the case of the linguistic meaning being more specific than the thought exists as well.

First, for the sake of clarity, I will give some brief comments on what vagueness/specificity is not.

One thing it is not is ambiguity. Pinker (1994) brings up examples like *Child's Stool Great for Use in Garden* (pp. 78-79) as indicative of inexactness in language. What is at play here is not just homophony—syntactic or constructional ambiguity has the same effect, as can be seen in the familiar joke *I shot an elephant in my pajamas*. This is essentially a matter of coincidence: two linguistic structures share the same form, and so it looks like there is inexactness in what the expression refers to, but it is inconceivable that these structures would not be distinguished in the mind of the language user. So, as far as we are concerned, this is a non-problem. But it does demonstrate the rather obvious point that the cognitive objects directly involved in

representations of linguistic meaning are the meaning-bearing elements of language and not the superficial externalizable forms attached to them (see Section 2.3 (The architecture of language) for more discussion).

Another issue from which it is to be distinguished—more subtlely this time—is context-sensitivity in what expressions refer to, as seen in, for example, deixis (c.f. Pinker 1994: 80; Levinson 1997: 19-20, 2003: 293-294). For instance, what *tomorrow* refers to depends on when it is said: if I say it today, it means tomorrow; if I said it yesterday, it means today; if I say it tomorrow, it means the day after tomorrow. *This* can refer to absolutely anything, depending on context: if one points at a cat, it can refer to a cat; if one points at the sky, it can refer to the sky; and anaphoric reference is just as freely variable depending on linguistic context. Likewise, *The serve was too fast to return* can be uttered about any number of different serves. This context-sensitivity of reference is partially similar to the case of ambiguity: *this* can refer to different things, but that does not mean that they are not distinguished in the mind of the language user. However, it is a different phenomenon in that the linguistic meaning does not change: the linguistic meaning of *this* does not depend on what one is pointing at; the linguistic meaning of *tomorrow* does not change depending on whether I say it today or yesterday. The issue here is of expressions being more or less *explicit*, but not more or less *specific*. This sort of 'looseness' is not the granularity we are concerned with.

Now we turn to vagueness/schematicity and specificity. Levinson (1997: 18) provides two examples of the sort of phenomenon we are interested in. The first is the fact that *some* can be taken to mean 'at least some' or 'some but not all', but can also have a general, schematic meaning that encapsulates both without having to specify between them. So *Some of my books are missing* (with the schematic meaning) does not have to commit to any the connotations that would attend the more specific meanings, like 'it is possible that all of the books are missing' and 'all of the books might have been missing, but only some are'. The second example is the schematicity seen in Guugu Timithirr conditional constructions: there is no word for 'if', and the construction used for conditionals is general over meanings like 'if A, then B' and 'possibly A and B' etc.

It may be tempting to assume that familiar categorical distinctions (e.g. 'if A, then B' versus 'possibly A and B') are absolute in specificity (and it is only expressions that do not make these distinctions that are vague), but for any particular linguistic meaning, it is generally possible to divide it into further-specific meanings.

Different levels of granularity are possible across expressions in general. Langacker (2003: 47/2009b: 6) gives the example of *thing* → *object* → *vehicle* → *truck* → *pick-up truck* → *battered old pick-up truck*; *thing* is very schematic, and *battered old pick-up truck* is relatively specific and detailed. This also applies to situations/events/scenes: Croft notes that the same event can be given a relatively coarse-grained description like *He takes the pears and he puts it on the bicycle* or a relatively fine-grained one like *And then he decides to take the whole bushel. And he puts it on his bicycle rack in front…* (2007: 356-357). Jackendoff (1996: 10-11/1997: 188) brings up the point that, for example, the expression *Bill killed Harry* can correspond to a number of specific scenarios—there are various means by which Bill might have killed Harry: stabbing, strangling, shooting, poisoning, etc. That is to say that *Bill killed Harry* is vague relative to any of those specific conceptions.

In Evans' examples *The teacher scrawled in red ink all over the pupil's homework exercise* and *The red squirrel is in danger of becoming extinct in the British Isles* (2009: 108/206), *red* refers to different colors. The meaning of *red* does not by itself specify from among these; in fact, it is possible to imagine an infinite variety of colors that are all 'red'. Indeed, variation in granularity of color terms across languages is a well-studied phenomenon. See, for example, Berlin and Kay (1969), Hardin and Maffi (eds.) (1997), Kay et al. (2009), Biggam (2012), and MacDonald, Biggam, and Paramei (eds.) (2018).

Tense also exhibits schematicity and (inter-language) gradience in granularity (Evans 2009: 112-113). Some languages make a binary distinction: something is either in the past or not in the past; exactly when it occur (beyond this distinction) is not relevant. Other languages with higher granularity make distinctions between recent past and remote past. Still higher granularity is exhibited by other languages, which have an exotically high number of distinct tenses. Evans describes this sort of schematicity as

arising from *parameterization*, whereby details are 'compressed' into a relatively few number of general parameters; this is essentially a "reductive abstraction" (p. 113).

The phenomenon of polysemy complicates this issue of schematicity in linguistic meaning, specifically because it has a convoluted relationship with schematicity. It is possible as a rule to at least pose the question *Is this schematicity/vagueness, or polysemy?* for any case. The import of this question for the present point is the issue of whether there is one (schematic) representation or several more specific representations. What we are concerned with is if the meanings are separately represented, with what degree of granularity the divisions are made; and where they are not distinct, what level of schematicity holds.

If the question is asked in the extreme, definitive answers may appear. For example, it would be untenable to posit polysemy for each of the infinite possible variations of colors that can be 'red'. However, to the question *Is there a degree of polysemy in* red*?*, the answer is less clear-cut: maybe. A single schematic meaning is in principle possible, but (non-infinite) polysemy is also possible: for instance, one meaning may cover colors that are actually some shade of red, and another denote things that are red-ish in hue but not actually red (like the red fox). And perhaps things that are figuratively 'red' may be serviced by separate meanings as well[38]. The point is that often the answer is not fixed; in many cases, nothing stops a language user from having a polysemic or schematic conception of a particular linguistic structure.

Another example is the well-known matter of the meaning of *in*. The physical arrangements referred to by *the water is in the bottle* and *the flower is in the vase* are distinguishable: the bottle encloses the water, but the flower sticks out of the vase (indeed, the point of a vase is that it sticks out). Evans (2009: 155-164) gives a polysemic analysis for *in*, arguing for two distinct meanings that he refers to as [ENCLOSURE] and [LOCATION WITH SURETY]. On the other hand, a non-polysemic—i.e. schematic—analysis is also possible, whereby differences in the particularities of the physical arrangement are all irrelevant because the object is *functionally* 'in' the thing

[38] Or not—c.f. the discussion of metaphor in Section 2.2.1.4.

(c.f. Jackendoff 2002: 354-355; see also Lakoff 1987: Ch.2 (pp. 416-461) for an analysis of various senses of *over*).

The fundamental problem is that there is no strict delineation between polysemy and schematicity. Often both are possible: for example, *love* can be polysemous between filial love and romantic love etc., or it can be schematic and non-distinguishing between different types of love; Langacker (2008: 224-226) brings up the similar example of *mail* being able to refer to specifically physical mail, specifically e-mail, or a schematic concept that does not distinguish between them. This applies to constructional meaning as well. One of the central problems in analysis of constructions is how schematic constructions and constructional meanings are (see, e.g., Goldberg 2006, 2009; Croft 2009; Langacker 2009a; Hayase 2018). This problem, as Langacker argues, does not have a general straightforward answer.

So granularity is a rather messy issue. Any particular expression has some degree of relative vagueness, and since it often is not possible to determine absolutely whether a particular linguistic element is in fact polysemous (and if it is, to what extent) or schematic, the actual specificity of linguistic meanings is not fixed (c.f. Tuggy 1993). This does not mean that any case of potential polysemy is also analyzable as schematicity. For example, the meanings of *declare* (at customs) and *declare* (war) are probably not validly analyzed as being instantiated by one schematic concept (c.f. Evans 2009: 232-234)—only a polysemic analysis is tenable. (Historically, situation-specific meanings can be entrenched and conventionalized as distinct meanings, effecting a transition from (definite) schematicity to (definite) polysemy. See footnote 37 above.)

My ‘answer’ to all of these questions about whether particular elements are polysemous or schematic and how polysemous they are when they are indeed polysemous is that it depends on the analysis of the language user—the sorts of possible variations discussed here are possibilities for what is actually represented in the head of the language user—and that is what matters. The issues here—whether, for instance, *some* or *in* or *love* is polysemic or schematic, and what specific meanings they are understood to express—just as discussed in Section 2.2.1.2 above, are matters of individual-dependent and situation-dependent variability in analysis.

Many researchers (e.g. Pinker 1994; Jackendoff 1996, 1997; Levinson 1997, 2003; Fodor 2001; Vicente and Martínez-Manrique 2005, 2008; Evans 2009) have remarked that linguistic representations are often less exact than the presumable underlying conceptual structure. Often, the specificity of the expression will be less than that of the intended thought. This is related to the general schematicity just discussed. If a language user describes a conceptualization using language, the linguistic expression will often not represent the full detail of the original conceptualization. For example, one may have a particular idea, and use the word *beauty* to refer to it, but 'beauty' does not denote the specific concept in the user's head; it is more general/schematic, so even if that idea can accurately be described as *beauty*, *beauty* does not represent the idea to the original specificity. The word *beauty* can refer to many different specific conceptions; its meaning is indeterminate between them.[39]

Thought in its language-independent conception can be of arbitrary detail. But the meanings available in particular languages are not so free. So while language has a degree of flexibility in that different levels of granularity of expression can be used for different levels of detail in conceptualization (as seen above (e.g. *thing* → *object* → *vehicle* → *truck* → *pick-up truck* → *battered old pick-up truck*)), the granularity of the expression will generally not match the granularity of the thought.

There are two cases of mismatch. One is the case where the thought is more specific than the best-available expression. If (to return to our previous example) there is no better word than *beauty*, no circumlocution will get the expression closer to the thought—and it often makes it worse, by adding unwanted information. Levinson (1997: 18-19, 2003: 292-293) points out that *Some of my books are missing* can potentially refer to books one owns, wrote, borrowed, etc.; however, if one attempts to specify the books from among these, for instance as *Some of the books I own*, the expression then may imply that the books one borrowed are not missing—which was not the original point, and so attempting to better capture the exactness of the thought in the expression has in fact led it further astray. The other case is that where the thought is vague relative

[39] The obverse of this, of course, is that schematicity in linguistic meanings affords language flexibility: a limited number of linguistic elements can thereby cover an indefinite number of more-specific meanings (c.f. Smith 2000: 348-351).

to the best available expression. One must not fall into the misconception that the thought must be precise where the expression is vague. If all one has in mind is 'pick-up truck', then *battered old pick-up truck* will, of course, be less faithful to the thought by virtue of being overly specific (even if it is, in fact, battered and old). There can be intended relative vagueness or non-distinctions in the thought. If in the original conception of the thought, there is no specification between books one owns, wrote, borrowed, etc., then any specification in the expression beyond *my books* will represent a deviation from the granularity of the thought. If the language user does not intend a distinction between singular and plural, but is obligated to express the concept in question with one of them or both, the expression is again too precise.

The issues discussed here show some aspects of the fact that linguistic expressions can generally do no better than approximate thought: generally either the thought is more exact than the expression, or the expression is more exact than the thought. See Section 4.2.1 (Categorization) for more discussion of the 'fit' between language and thought.

2.2.1.4 The content of linguistic meaning

In this section, we will discuss aspects of the content of linguistic expressions. We will address four topics: propositions, imagery, metaphor, and structural relationships. These are, of course, not at all intended to exhaustively characterize linguistic meaning, but they do constitute important aspects of it.

2.2.1.4.1 Propositions

The content of linguistic expressions is very often described as propositional. Propositions are 'about' things in the same sense that linguistic sentences are 'about' things. Propositions also have the same sort of subject-predicate or verb-argument structure that is a fundamental part of linguistic sentences. As common as propositionality as a characterization of the content of expressions is, however, it is problematic for a few reasons.

The first is that it does not have a consistent meaning across the literature. There are two main conceptions of "propositionality", that are contradictory.

One takes propositions to be abstract, (linguistic-)representation-invariant objects, so that two (or more) different sentences can express the same proposition. Of course, two different sentences will necessarily have different meanings, so this notion of proposition corresponds to the 'gist' of the meaning expressed by sentences. In this sense, for instance, both *Marshall gave Rick a watch* and *Rick received a watch from Marshall* can be said to express the same proposition (Barsalou 1999: 595). People often remember the gist of expressions without remembering the specifics, and this sort of "reduced" proposition seems to correspond to that content. Now, this sort of object falls quite short of characterizing the content of expressions: it carries a loss in information compared to the actual meaning of expressions. The gist is derived from the meaning of the expression, but it is not the meaning itself[40].

The other conception of propositions takes the opposite stance, and characterizes them as construal-dependent objects. So, depending on the construal, a particular situation can be propositionalized (for example) as *ABOVE (ceiling, floor)* or *BELOW (floor, ceiling)* (Barsalou 1999: 595[41]). Here, different sentences necessarily express different propositions. This conception is certainly much better at characterizing the content of linguistic expressions.

It is interesting to note that the same phenomenon of remembering only the gist has been the basis for two opposite conclusions concerning the format of the relevant mental representation (and also the format of the original thought to be expressed in language). One, formed on the basis of the first conception of propositions described above, posits that it is propositional (c.f. Levelt 1989: 74-45; Barsalou 1999: 595). The

[40] Gist mental representations may, however, be relevant in discourse contexts or 'long' thoughts where working memory would otherwise be overloaded. Presumably, either the more detailed representation is transferred to a gist representation or they are formed in parallel and the detailed representation is later lost.

[41] One may note that it is odd that Barsalou discusses both of these conceptions of propositions—in immediate succession—as being valid without remarking on the contradiction.

other, formed on the basis of the second conception of propositions, describes it as non-propositional (Chafe 1977; Croft 2007).

The second reason propositionality is problematic as a characterization is that it can imply different things that are often not part of the actual representation.

One of these is truth values. Though propositions, by virtue of the fact that they can be true or false, conveniently facilitate a formal description of meaning in terms of truth conditions[42], truth values are totally emergent and secondary: they do not constitute meaning. As Jackendoff (2002: 326) notes, in, for instance, *The apple on the table astounded Max* or *That the Red Sox won today astounded Max*, what astounded Max was not a truth value, but an event. Truth values are at most a theoretical tool for describing meaning. They are not actually part of the meaning of expressions: unless the expression is specifically about something being true or false, truth is entirely secondary; and even then, it must be noted that various separate linguistically-encoded concepts can be related to truth (e.g. *It is true that…*, *It is certain that…*, *It is the case that…*, *Factually,…*), and these are not identical in their meaning or—it then goes without saying—their mental representation.

Only in some cases is what I will describe as *propositional force*—a connection to whether that proposition is the case or not—even relevant at all to the meaning of linguistic expressions. One can assert, deny, or entertain a proposition; in any case, whether it is the case or not is a central part of the consideration. For instance, central to the meaning of *If A, then B* is whether *A* and *B* are the case, and the overall expression can be paraphrased as 'If it is the case that A, then it will be the case that B'. When linguistic expressions involve epistemic modality, they deal with propositional force and propositional attitude (i.e. the epistemic stance with respect to some entertained proposition). For example, *He must have eaten the apple* and *I doubt he ate the apple* express propositional attitudes toward the proposition that he ate the apple. But these are special cases. Propositional force is not fundamental to linguistic meaning in general. First of all, propositional force cannot exist for expressions that do not reach the size of proposition (these do not have to be fully-explicit propositions—c.f. Polo-Sherk 2016a,

[42] The truth conditions for a proposition are simply the conditions or set of situations under which a proposition would be true.

2016b) like the example of *beauty* referred to earlier in 2.1 (though a single word like this can in some cases instantiate a proposition (e.g. as in *Its beauty was unexpected*, which can be paraphrasable as 'The fact that it was beautiful was unexpected'). Propositional force is also irrelevant in many cases. To take a sentence from the text above, *Next, we will consider two separate but related issues*, while understandable if rephrased as something like 'It is the case that we will next consider two separate but related issues', is not really about an assertion of that proposition. Most expressions, even if they are consistent with some propositional attitude or truth value, do not actually have propositional force or truth as part of their meaning.

It should be noted that there is a further, narrower conception of proposition that limits propositions to certain kinds of semantic content: for example, the underlined bits in *The idea that mangos are inferior to kiwis is preposterous* or *I want for you to be happy* qualify as propositions in this sense. In terms of linguistic structure, these generally correspond to complements, but not necessarily entire sentences.

While a large part of linguistic meaning can be described as propositional, propositionality does not generally *explain* linguistic meaning. Levelt (1989), in adopting propositionality as a characterization of linguistic meaning, notes "[W]hen something is predicated about a referent, that predication can be true, or false, or undecidable. Such representations are often called *propositional.* Let us adopt this practice, but be aware that messages do not always have truth values. When you say *Congratulations!* or *What?*, the underlying message is probably not a proposition. So, "propositional" will stand for a mode of representation of which propositions are a special case." (pp. 70-71). To the extent that propositionality is accurate as an overall descriptor of the content of linguistic meaning, it simply means 'linguisticness'. And so it is not very helpful in the end.

2.2.1.4.2 Imagery

Another aspect of the content of the meaning of expressions is imagery (which is not necessarily visual or even spatial). Some parts of conceptualized imagery can be

depicted in linguistic expressions. This is something that has been researched quite a bit in cognitive linguistics, and so there are ways to talk about it.

One part of imagery is perspective. For example, a particular physical arrangement can be described alternatively as *the ball is in front of you* or *the ball is to the right of the lamp* (Levelt 1989: 51-52), this reflects differences in perspective (particularly, differences in what is taken as the reference point, as mentioned in Section 2.2.4.1). Also, different vantage points can be assumed, as seen in, for instance, the contrast between *I'll go to your apartment tomorrow* and *I'll come to your apartment tomorrow* (Langacker 2003a: 252). Additionally, perspectives can be *event-external*, as in *I like Chicago*, or *event-internal*, as in *I like it in Chicago* (Machida 2017)[43].

Another is the mental scanning (Langacker 1987, 1997, 1999, 2000, 2008, 2009b) mentioned earlier. *The roof slopes steeply downward* and *The roof slopes steeply upward* reflect dynamic images (Langacker 1997: 244). Mental scanning can relate to more abstract, non-spatial imagery as well. For example, *Don't mention calculus—elementary algebra is already too advanced for him* (Langacker 2008: 82) and *From the brightest to the dumbest, the students all work very hard* (Langacker 2008: 534) exhibit mental scanning in non-spatial abstract imagery. Other variations of this phenomenon include what Langacker calls *fictive/virtual change* and *fictive/virtual motion*. Fictive motion is the use of movement-imagery to describe a static situation, as in *That mountain range goes from Mexico to Canada.* (Langacker 1999: 82) or *There is a house every now and then through the valley.* (Talmy 2000a: 71). The mountain range is not

[43] The general situating of linguistic content in the discourse context is known as *grounding* (Langacker 2008: 259-309); this can be achieved through determination (*a*/*the*/*this*/etc.), for example. Now, some languages do not use determination (such as Luiseño (a Californian Uto-Aztecan language); e.g. *hunwut xaari-q* (lit. 'bear growl'; '{The/A} bear is growling') (Langacker 2008: 496), and some languages can (additionally) omit the subject (such as Japanese; e.g. *waa, butsukaru!* (lit. 'whoa! (going to) hit!'; 'Whoa! We are gonna hit 'em'); here, unlike in English, specifying the subject and object would be quite odd (c.f. *#waa, watashi-tachi ga mae no kuruma ni butsukaru!* (lit 'Whoa! We (are going to) hit (the) car in front!') (Machida 2020: 249-251)). Machida (2012, 2013, 2017, 2020) proposes that the (covert) grounding in these cases is effected by portraying the content in an event-internal manner.

actually going anywhere; the houses are in fact in constant existence, but if one imagines moving through the valley, they can be described as being there 'every now and then'. Fictive change is the depiction of a static correlation as a changing relationship, as in *Our Christmas tree gets smaller every year.* (Langacker 2008: 530) or *As body size increases, the average gestation period gets longer.* (ibid.: 534). There are actually several Christmas trees, none of which actually changes size. The body size of any particular animal is not increasing at all; rather, the body sizes and gestational periods are mentally scanned over the set of animals, which results in a dynamic (i.e. "changing") construal of the correlation relationship between them.

There are also what are called *image schemas* (Johnson 1987; Lakoff 1987; Gibbs and Colston 1995; Grady 2005; Mandler and Pagán Cánovas 2014; etc.)—abstracted structural bits of images that are necessary for imagistic conceptualization; these form the structure of an interpretable image. Johnson proposes image schemas such as CONTAINER, BLOCKAGE, CONTAINMENT, ENABLEMENT, CYCLE, and PART-WHOLE as basic to imagistic conceptualization and understanding. Mandler and Pagán Cánovas (2014) discuss the role of image schemas in the development of the ability to interpret physical scenarios in infants: for example, infants first perceive motion into and out of containers more than the containers themselves, and therefore tend to make the overgeneralization that an object that has moved behind a screen will be blocked from view. Image schemas can be reflected in linguistic expressions. For instance, there are different ways of talking about ways in which something is 'in' something else, or in something else's way, or in a part-whole relationship; dynamic image schemas like CYCLE can likewise be depicted (e.g. *read a book for an hour* is not cyclic, but *jump up and down for an hour* is); abstractions are also possible (e.g. *All of the glory went to me* reflects the use of a path-goal schema in an abstract context). So image schemas are part of imagistic conceptualization, and they can form an abstract image that can be described in linguistic expressions.

Now, while much of the conceptual structure reflected in linguistic expressions can be described as imagistic in some sense, not everything related to conceptualization or construal is imagistic. For example, perceiving something as ugly or beautiful, perceiving something as something that is closer to contempt or to disdain, or thinking

about the government's justification for war as *exaggerated* or merely *flawed* (October 6, 2004 BBC radio interview, reported in Croft 2007: 345) are matters of construal but not really of imagery.

2.2.1.4.3 Metaphor

Another extensively-studied aspect is metaphor.

Metaphorical expressions can reflect a conceptual mapping. In cognitive linguistics, metaphor is described as a mapping from a *source domain* to a *target domain*. For example, in the TIME IS SPACE metaphor, conceptual material from the domain of space is mapped to the domain of time; this mapping is realized in a great many expressions across languages. Moore (2014) identifies three basic patterns: MOVING EGO (e.g. as in *We're headed for fall. Spring is behind us*. (p. 8)), EGO-CENTERED MOVING TIME (as in *Summer is coming. Winter is gone.* (p. 12)), and SEQUENCE IS RELATIVE POSITION ON A PATH (as in *Spring follows winter.* (p. 66)). In all of these cases, conceptual structure concerning spatial movement and position is applied to time. In moving ego metaphors, time is stationary and the ego moves with respect to it, from past to future. In ego-centered moving time metaphors, the ego is stationary and time moves with respect to it, from future to past.[44] In sequence is relative position on a path metaphors, the ego is irrelevant, and events are spatially ordered as if they were moving toward their realization, such that earlier events—in the relative past—are ahead of later events—in the relative future (e.g. spring is in the future relative to winter). Metaphorical space-to-time mappings can be seen in gesture, which can be in accord with linguistically-established mapping patterns (e.g. gesturing backward to indicate the past; Casasanto and Jasmin 2012; Casasanto

[44] This wording may be confusing, due to the polysemy of the word *time* (c.f. Evans 2005), to those unfamiliar with research on temporal metaphor. For example, one may speak of a *river of time*, where time flows toward the future (i.e. the opposite direction of ego-centered moving time). The *river of time* conceptualization is essentially an egoless moving ego: in moving ego, the point in time occupied by the ego moves, and therefore 'time' in the sense of the flowing river—but not in the sense of ego-centered moving time—moves (like the river) from past to future.

2016; Núñez and Sweetser 2006[45])—including in languages that use absolute frames of reference (Gaby and Sweetser 2017; see Levinson 2003: 244-271 for a general discussion of gesture in languages with absolute frames of reference)—or not (e.g. gesturing on a left-right 'timeline' despite the absence of left-right linguistic temporal metaphors; Casasanto and Jasmin 2012; Casasanto 2016).

Conceptual metaphor theory (Lakoff and Johnson 1980; Lakoff 1987, 1993; Grady et al. 1996; Gibbs 2011; Kövecses 2015; etc.) posits that the target concept is understood in terms of the source concept. For example, in the TIME IS SPACE metaphor, time is understood in terms of space. Some versions (e.g. Lakoff and Johnson 1999; Gallese and Lakoff 2005; Johnson 1987, 2007; Lakoff 2008; Kövecses 2015) emphasize an embodied experiential basis to conceptual structure in general. The idea is that concrete concepts are understood through sensorimotor experience, and their conceptual structure mapped to more abstract domains to enable understanding of abstract concepts. This means that mental representations of metaphorical expressions would represent the "literal/concrete" meaning (e.g. physical warmth in *a warm welcome* or physical movement in *summer is coming*); and at a further extreme, all conceptual representations would involve sensorimotor representations.[46]

Conceptual metaphor theory has inspired a significant amount of experimental research (see Gibbs 2011 and Minervino et al. 2018 for reviews). For example, McGlone and Harding (1998) found that the interpretation of the ambiguous metaphorical statement *The meeting originally scheduled for next Wednesday has been moved forward two days* (which can be taken to mean either that the meeting is now on Monday, or that it is now on Friday) can be affected by priming with either moving-ego-framed expressions (e.g. *We are approaching the deadline*) or moving-time-framed expressions (e.g. *The deadline is approaching us*), such that *forward* tends to be interpreted either in accordance with the moving-ego perspective (i.e. toward the future: Friday) or the moving-time perspective (i.e. toward the past: Monday), respectively.

[45] Aymara, studied by Núñez and Sweetser, is famously exceptional in that it depicts the ego as facing the past (e.g. 'ahead' refers to the past, rather than the future (as in English and many other languages)). Gesture in Aymara reflects this.

[46] See Section 2.2.2 (Modal versus amodal representation) for more discussion of embodiment.

Wilson and Gibbs (2007) found that, for example, performing the physical act of grasping facilitates the comprehension of linguistic expressions like *grasp the idea*.

Neurological studies on conceptual metaphor theory have produced mixed evidence. For example, Boulenger et al. (2009) found activation in brain motor areas related to upper and lower limbs in comprehension of the expressions *He grasped the idea* and *He kicked the habit* (importantly, these are established, conventional idiomatic expressions, and not novel ones); Desai et al. (2011) obtained similar results. However, Aziz-Zadeh et al. (2006) and Desai et al. (2013) had negative results.

Strong conceptual metaphor theory has been repeatedly criticized (e.g. Haser 2005; McGlone 2007, 2011; Madsen 2016). The strong version of conceptual metaphor theory suffers from two related critical problems: the assumption the metaphorical expressions necessarily reflect concrete-to-abstract conceptual mapping and the rejection of conceptual abstraction.

It assumes all linguistic expressions that can be theoretically analyzed as metaphors are actually active as metaphors cognitively. This is problematic partially for reasons related to analyzability and historical change: for example, *before* is basically exclusively temporal in modern English (irrespective of its historical spatial meaning); *precede* (formally <before> + <fall>) is likewise non-spatial (it is either temporal or abstract, in the sense of 'priority'); *precede* is probably not analyzable for many people, in which case any latent spatial sense is unextractable, and at any rate, even if it is analyzed as 'fall before', there is no reason *pre* ('before') must be analyzed spatially and not temporally or abstractly. Binder and Desai (2011) suggest that only novel metaphors are understood through sensorimotor representations, and established ones are understood directly and abstractly (though note the Boulenger et al. results mentioned earlier).

It also denies abstraction. For example, in strong conceptual metaphor theory, it is asserted that time is abstract and therefore can only be understood through conceptualizations of space. This has been criticized by Evans (2013), who argues that time can be understood in its own terms, and McGlone (2011: 568-569), who also notes that his own earlier findings (McGlone and Harding 1998) do not warrant the

conclusion that space is actually understood in terms of time.[47] To give another example, the conduit metaphor (c.f. Reddy 1979; Lakoff and Johnson 1980; Grady 1998), which analogizes communication and transfer (as in *Try to get your thoughts across better* (Reddy 1979: 286)) is assumed to be mentally instantiated by a mapping from spatial movement to informational transfer; but there is no reason informational transfer cannot be understood directly, or a through a single abstract concept of 'movement' that subsumes both spatial physical movement and informational movement.

These problems affect the validity only of the strong versions of conceptual metaphor theory. Weaker versions (e.g. Gibbs 2011) that do not assume universality and necessity of literal/concrete-to-abstract conceptual mapping are certainly valuable; analysis of metaphor can indeed provide insight on conceptual structure—it just should not be assumed that everything is metaphor cognitively.

Gibbs (2009: 22-23) notes that both supporters and detractors of conceptual metaphor theory may be basing their assumptions on their own introspection; but if we assume that both introspections are right (i.e. that some people understand particular concepts metaphorically and other people do so abstractly), this is problematic only for the strong conceptual metaphor theory, which asserts that the mental representation *must* be metaphorical and not abstract. I am not aware of any studies investigating the possibility of individual differences in the processing of metaphorical language. Although it is generally difficult to draw conclusions from fMRI data at the individual level (because patterns of brain activity vary very greatly between individuals), individual differences are detectable in other contexts: for example, Miller et al. (2012) found that individual tendencies to utilize visualization or verbalization in memory recall correlate with differences in brain activation patterns observed with fMRI; see also Parasuraman and Jiang (2012). Individual differences may also be a factor in the

[47] Also, the "unidirectionality" of many metaphors (e.g. time is frequently talked about in terms of space, but space is generally not talked about in terms of time) has been cited as evidence for the concrete-to-abstract mapping structure of conceptual systems posited by (in particular strong versions of) conceptual metaphor theory. As McGlone (2011: 569) points out, this is not evidence that the relevant concepts are actually metaphorically structured: ""An insult is a razor" is more apt and meaningful than "A razor is an insult," but that in and of itself is not evidence that the conceptual structure of insults is predicated on razor knowledge."

larger issue of the role of sensorimotor information in the mental representation of concepts discussed in Section 2.2.2 below.

Finally, it should also be stressed that even if the mental representation of some expression is not actually metaphorical, it still reflects conceptualization, just like any other aspect of meaning.

2.2.1.4.4 Structural relationships

Another important aspect of conceptual structure depicted in linguistic expressions is what I will call the *structural relationships* in a complex expression (i.e. one that is comprised of several parts). This refers to structural configuration of the various conceptual objects—how the different parts are conceptually structurally related to each other.

One generally simple and clear example is parallelism: different parts can have separate but equivalent conceptual roles, or else some other non-parallel relationship.

In *The move was over-planned and poorly-executed*, *over-planned* and *poorly-executed* are separate and parallel properties of the move. There are different linguistic structures that correspond to different structural relationships. *The ball was hard to see and difficult to catch* is ambiguous here. One of the possible structural relationships is that in which *hard to see* and *difficult to catch* are just separate properties of the ball—they do not interact with each other, and are parallel. Another one, made possible by a different <and> construction (which is not necessarily overtly differentiated from the parallel one), is that being difficult to catch is a consequence of it being hard to see (i.e. there is a cause-effect relation); here, they are not parallel. This sort of parallelism and non-parallelism is related to the syntactic notions of *symmetric* and *asymmetric coordination*. In the following example, translated from the Brazilian language Kĩsêdjê, the bolded clauses are symmetrically coordinated (i.e. parallel) and the underlined clause is asymmetrically coordinated (i.e. non-parallel) and here expresses a cause-

effect relation. ***At the school there are no mosquitoes*** *and* ***my shirt was dirty*** *and (so) I didn't put it on* (Nonato 2014: 117[48]).

The kinds of parallelism that are expressible depend on what constructions are available. Different languages permit the expression of different kinds of parallelism.

For example, in *That is and always will be a mistake*, there is parallelism between [is] and [always will be]. It is not possible to preserve this parallelism in Japanese because Japanese does not have an equivalent construction[49]; the closest expression would be something like *sore wa machigai de ari, itsumo machigai de aru*, which has a structure like *That is a mistake and always will be a mistake*, in that the parallelism only inheres at a higher hierarchical level, between [is a mistake] and [will always be a mistake].

Nominative-accusative systems and ergative-absolutive systems enable fundamentally different patterns of parallelism. In a nominative-accusative system, agents of verbs (whether transitive or intransitive) are treated alike (as nominatives), and contrasted with patients of verbs (treated as accusatives); in an ergative-absolutive system, agents of transitive verbs (treated as ergatives) are contrasted with agents of intransitive verbs and patients of transitive verbs (treated as absolutives). Comrie (1981: 105-109) points out the contrast in parallelism that this results in with the following example.

i) a. The man hit the woman. (ibid.:106) (English)
   b. The man came here.
   c. The woman came here.
   d. The man hit the woman and came here (= (a) + (b))

[48] Nonato gives *…and (then) I didn't put it on*. I have used *so* instead here simply because it is more idiomatic in English. In other situations, *and then* in English can, of course, express similar cause-effect relations (e.g. *If I were king of the world, I would give everyone a cookie, and then everyone would be happy.*).

[49] Specifically, Japanese does not allow coordination of tensed verb phrases, and additionally does not distinguish between present tense and future tense, using a single "non-past" tense that covers both.

ii) a. Balan-d^{y}ugumbil baŋgul-y aŋgu balgan. (ibid.) (Dyirbal)
woman-ABS man-ERG hit
'The man hit the woman.'

b. Bayi-y a banin^{y}u.
man-ABS came.here
'The man came here.'

c. Balan-d^{y}ugumbil banin^{y}u.
woman-ABS came.here
'The woman came here.'

d. Balan-d^{y}ugumbil baŋgul-ya aŋgu balgan, banin^{y}u. (= (a) + (c))
'The man hit the woman, and the woman came here.'

In English (a nominative-accusative language), the relationships between *The man* and the verb phrase (*hit the woman* and *come here*) are analogous; and so in (i-d), [hit the woman] and [came here] are parallel with respect to [the man]. (I should stress for clarity that they are asymmetrically coordinated and are not parallel with respect to each other. The linguistic phenomenon here is not quite the same as the one discussed above, concerning parallelism in symmetric coordination, but it is still an instance of the general phenomenon of parallelism.) However, in Dyirbal (an ergative-absolutive language), this is not so. The relationship between the Dyirbal equivalents of [the man] and [hit the woman] is different from that between the equivalents of [the man] and [came here]. In (ii-d), [hit the woman] and [came here] are *not* parallel with respect to [the man]. In fact, rather, [the man hit] and [came here] are parallel with respect to [the woman] ([came here] takes as its subject the absolutive, which is [the woman]). Essentially, the superficial structure *The man hit the woman and came here* means 'The man hit the woman and the man came here' in a nominative-accusative system, and 'The man hit the woman and the woman came here' in an ergative-absolutive system.

Zeugma is an example of a 'forced' pseudo-parallelism. The parallelism fails to establish fully because, due to the senses of the central element being different, the relation it has with respect to each would-be-parallel element is different. This can be seen in the following examples.

i) a. I saw a bunny and that Mama was right.
   b. He attacked me in a rage and a golf cart.
ii) a. I saw a bunny and saw that Mama was right.
   b. He attacked me in a rage (and) in a golf cart.

The pseudo-parallelism in zeugma depends on the various senses—here, of *saw* and *in*—sharing the same form (which usually is not a coincidence, as zeugma generally relies on semantic extension) so they can collapse into a single syntactic element. The fine structures of each sense, however, must simultaneously exist (in an overlaid fashion) in order for the meaning to be created. This can be contrasted to the 'de-zeugmatized' examples in (ii), where the overlaid structures are separated out. In these cases, parallelism is generally not destroyed completely, but it retreats to a higher hierarchical level: e.g. (i-a) has pseudo-parallelism between [a bunny] and [that Mama was right] with respect to [saw]; (ii-a) has parallelism between the full propositions [saw a bunny] and [saw that Mama was right], which are parallel with respect to [I]. Langacker (2008: 408) notes that while the parallelism in *She signed the papers [reluctantly] and [with much hesitation]* is well-formed, it is perhaps not so in *She signed the papers [reluctantly] and [with a ballpoint pen]*, and is clearly not in *She signed the papers with [reluctance] and [a ballpoint pen]*, because the semantic roles of [reluctance] and [a ballpoint pen] are not equivalent.[50]

Another arena where structural relationships are particularly salient is the conceptual structure describable in terms of *mental spaces* (Fauconnier 1985, 1997, 2014; etc.). Mental spaces are abstract conceptual "domains" [51] ; linguistic representations can organize conceptual material according to these domains and relate material in different domains to each other in various ways.

[50] A variation of zeugmatic overlaying of structures that does not really exhibit parallelism can be seen in *Send your girlfriend somewhere really cool, the fridge for a pork pie* (alcohol advertisement, reported in Evans 2009: 230).

[51] The notion of *domain* in mental space theory is different from that in metaphor theory, in that it is an organizational space rather than a thematic one, but it does subsume it (i.e. metaphor can be described as mental space mapping; c.f. *conceptual blending* (Fauconnier and Turner 1996, 2002; Fauconnier 1997, 2014; Grady et al. 1999; etc.)).

In, for example, *She thinks that he is ambitious*, [he is ambitious] is not directly about reality; rather, it can be said to exist in a "belief space".

Conditionals can also be described as involving mental spaces (for an extended analysis of conditionals in terms of mental spaces, see Dancygier and Sweetser 2005). Conditionals can be described as setting up a mental space containing the conditional protasis and then another mental space containing the content that is contingent on the conditional (the apodosis). For example, in *If I'm hungry, I'll eat an apple*, firstly, [I'm hungry] can be said to belong to an "if space". The material contingent on that (i.e. [I'll eat an apple]) can be said to belong to an extension of that space.

Material in various spaces can relate to each other in more complex ways. For an example, we can consider the conceptual structure expressed by *If I were my boss, I would fire me* in terms of mental spaces. In reality (i.e. the "base space"), [I] and [my boss] exist (separately). Here, the imagined hypothetical space set up by *If I were my boss* identifies [I] with the *role* of [my boss] (that is: not the particular person who is actually my boss; c.f. *My boss said I did a good job*, where it is the actual boss—the *value*—who said I did a good job). In the extension of this imaginary space, I-as-my-boss can fire I-as-myself (i.e. the actual me)[52].

Fauconnier (1997: 44-48) illustrates the function of interactions between spaces with the following example: *Achilles sees a tortoise. He chases it. He thinks that the tortoise is slow and that he will catch it. But it is fast. If the tortoise had been slow, Achilles would have caught it. Maybe the tortoise is really a hare.* Here, the first two sentences are about material in the base space. The third contains material in Achilles' belief space; there are two layers of this material: [the tortoise is slow] and [he will catch it], which is consequent of the former[53]. The fourth sentence returns to the base space, and is linked conceptually with [the tortoise is slow] in Achilles' belief space. The fifth sentence has a hypothetical space, containing the protasis [If the tortoise had

[52] Note also that here *I would fire me* is possible; normally, it would have to be *I would fire myself*.

[53] As mentioned above, another possibility is that [the tortoise is slow] and [he will catch it] are parallel and not in a cause-effect relationship. Also, as Fauconnier says, the future tense of *he will catch it* can mean that the conceptual material belongs to a "future space".

been slow], which is linked to [the tortoise is slow] and [it is fast] in previous sentences, and (in an extension) the apodosis [Achilles would have caught it], which is linked to [he will catch it] in the third sentence[54]. The sixth sentence returns to the base space, and sets up a possibility space in which [the tortoise] (linked, of course, to all previous references to it) is identified with [a hare].

Conceptual material belonging to different mental spaces can also overlap in the linguistic structure they correspond to. For example, in *Jan believes that her son is **smarter** than he really is*. (Langacker 1997: 15), the conceptual material corresponding to the underlined section pertains to Jan's belief, and that corresponding to the bolded section pertains to reality (or, at least, they pertain to what the speaker believes to be Jan's beliefs and reality). Here, the two sets of conceptual material are that there is some degree of smart that Jan believes her son is, and that that believed degree of smart is greater than the actual degree of smart that he is. In the structure of the linguistic expression, however, they are not separated out, but rather overlap.

These are only a few aspects of the conceptual organization and structural relations expressed by language. Structural relations are an important part of conceptual structure in general, and they are another part of the content and meaning of linguistic expressions.

2.2.2 Modal versus amodal representation

A subject that is currently under debate within cognitive science is whether the mental representation of concepts has a perceptual base (or a *modal* format) or not (or an *amodal* format). (The thesis of the perceptual representation of concepts is also

[54] It should also be noted that this hypothetical space is (probably) based on/subordinate to the base space; this can be contrasted with something like *Achilles thought that if the tortoise were slow, he would have caught it*., where the hypothetical space is instead based on the belief space. Rhetorically, it is possible that the fifth sentence is, in fact, intended to be from Achilles' point of view; in this case, the hypothetical space would be based on the belief space. The same applies to the sixth sentence.

sometimes referred to as the *embodied* theory of representation or (particularly in philosophy) as *(neo-)empiricism*.)

The modal hypothesis (Barsalou 1999, 2008; Prinz 2002; Glenberg and Kaschak 2002; Barsalou et al. 2003; Gallese and Lakoff 2005; Glenberg and Gallese 2012; Bergen 2012; Kaschak et al. 2014; etc.) posits that experienced perceptual states (this includes external-sensory and proprioceptive perception, motor representations, and also introspective perception (e.g. emotions)) are re-enacted in a context-dependent manner (i.e. only relevant information is activated) to mentally represent concepts. This is known as *simulation*, and will be addressed again in Section 3.3.3; here, I am only discussing the representational format itself. Simulation generally does not involve a total, exact reproduction of earlier perceptual representations, but rather is based on representations of perceptual schemata abstracted from experience. The representation of a concept is achieved by a multi-modal integration of modality-specific sensory information.

A somewhat separate motivation for the modal view comes from an extremely strong version of conceptual metaphor theory (e.g. Lakoff and Johnson 1999; Gallese and Lakoff 2005; Johnson 2007), which asserts that all abstract concepts are derived through metaphor from concrete sensorimotor experience[55].

It has been clearly established empirically that sensorimotor cognitive processes are related to conceptual processing (for reviews, see Chaterjee 2010; Kiefer and Pulvermüller 2012; Pecher 2013; Hauk and Tschentscher 2013; Leshinskaya and Caramazza 2014; Mahon and Hickok 2016; Galetzka 2017; Montefinese 2019). Behavioral experiments have shown interaction between conceptual processing and perceptual processing or perceptual information. For example, Glenberg and Kaschak (2002) found that button-push responses were faster when the direction of pushing matched the direction of action in stimulus sentences (e.g. toward the body for *Andy*

[55] A weaker version would postulate that some aspects of some concepts for some people may be informed by a metaphorical connection with other concepts. As discussed above, while I reject the strong version, the weak version is possible. It should be noted that Barsalou rejects strong conceptual metaphor theory (1999: 600).

gave you the pizza and away from the body for *You gave Andy the pizza*). Wheeler and Bergen (2010) had participants press a button to indicate whether a stimulus sentence made sense or not with their hand in a pre-assigned shape (either an open palm or a fist), and found that responses were quicker when the action depicted in the sentence matched the hand-shape (e.g. palm for *The waiter is smoothing the tablecloth* or fist for *The lawyer is carrying the briefcase*). Also, activation in sensorimotor areas of the brain has been shown to occur in the context of conceptual processing. For example, Hauk et al. (2004) found exposing participants to words pertaining to actions of particular body parts prompted activation of the relevant motor areas in the brain (e.g. *kick* would elicit activation in the leg area) (c.f. the later Boulenger et al. (2009) results mentioned earlier). Gonzalez et al. (2006) similarly found activation in olfactory areas elicited by words associated with odor, and Kiefer et al. (2008) found activation in auditory areas for words associated with sound. Pulvermüller et al. (2005) found that magnetically stimulating (transcranial magnetic stimulation) the relevant motor areas of the brain sped up processing of action words (e.g. stimulating the leg area resulted in faster processing of the word *kick*).

There are two main objections that can be raised against the modal hypothesis.

One major objection is that sensorimotor information (or sensorimotor information alone) may not be able to instantiate representations of concepts. An immediately obvious question would be how abstract concepts could be represented through perceptual information (c.f., e.g., Patterson et al. 2007; Shallice and Cooper 2013; Dove 2016, particularly pp. 1114-1115; Pecher 2018).

Barsalou (1999) demonstrates how some abstraction could be achieved in perceptual representational systems: for example, the concept of [truth] can be generated by comparing one's perception of reality with a perceptual simulation of some situation that is being entertained as possibly true; the schema abstracted from repeated successful mappings between these two representations is the representation of [truth]. However, in commentary on that article, Adams and Campbell (1999), Mitchell and Clement (1999), and Ohlsson (1999) point out that this may generate something of a concept of matching, but not exactly of truth; Barsalou's proposal may be correct as a description of a procedure for *evaluating* truth, but perhaps not as a mechanism for

representing the concept of truth. It should be noted that at any rate, this matching is abstract and emergent from identities between representations, and does not depend on a perceptual format of mental representation.

Additionally, there may be neurological differences in the mental representation (or at least in the processing) of abstract and concrete concepts. Different brain regions have been found to be activated in the processing of concrete and abstract words: the left inferior frontal gyrus, in particular, has been found to be associated with processing of specifically abstract words (Wang et al. 2010; Shallice and Cooper 2013; Dove 2016; Della Rossa et al. 2018; Montefinese 2019).

However, there is reason to be cautious about hypothesizing that the representational formats of abstract and concrete concepts are fundamentally different: there is no discrete distinction between 'abstract' and 'concrete'. For instance, [drill] (as in training) refers to physical action, but it is also 'abstract'. To give another example, 'beautiful' is sort of abstract in that it is quite schematic (e.g. compared to 'red'). It can refer to "concrete" beauty (e.g. *a beautiful sunset*), or abstract beauty (e.g. *a beautiful mathematical equation*). The beauty of a mathematical equation may be quite abstract, but at least it is plausible that it is generalized from a more concrete concept of 'beauty'. In this sense, other concepts are perhaps more abstract (e.g. 'mental representation' does not have a connection to concrete concepts).[56]

Shallice and Cooper (2013) (correctly) identify problems with conflating different sorts of abstractness, and suggest that imageability be used as the defining criterion. However, while imageability may indeed turn out to be the factor differentiating activation of the left inferior frontal gyrus (and potentially other brain regions), as a

[56] Kiefer and Pulvermüller (2012) make a similar argument (pp. 820-821). However, they endorse the position that sensorimotor representations can cover both concrete and abstract concepts. For example, "[f]or action-related abstract concepts such as "to free"—which can refer to a wide range of basic actions associated with freeing such as the removal of constraints of some sort—action representations related to these different kinds of freeing actions are obviously most crucial and should recruit motor circuits." (p. 820). I would point out that, conceptually, freeing from abstract constraints does not seem to involve physical action at all, so unless this is mentally represented as a metaphor from physical freeing, it is not clear why motor circuits should be involved in principle.

theoretical matter, it does not seem to definitively characterize abstractness in general: imagery can involve abstract concepts (for example, the subjective experience of beholding a beautiful equation can be imaged), and imagery itself can be abstract to a degree (e.g. as in *elementary algebra is already too advanced for him* (Langacker 2008: 82)).

Dove (2016) argues for a distinction between *generality*, *flexibility*, and *disembodiment*. On generality, Dove explains: "A concept such as MAMMAL may sit at the top of an abstraction hierarchy, and is thus more abstract than related lower-order concepts, but its referents are nevertheless concrete, perceivable objects. A concept like ODD NUMBER, on the other hand, is abstract because its referents are not concrete, perceivable objects but, instead, something more ephemeral." (p. 1111). Flexibility refers to the context-dependence of concepts and schematicity of the content mentally represented; "concrete" words can be used abstractly in metaphor, for instance. Disembodiedment is the conceptual divorcement from perception, for example as seen in concepts like [odd number], [justice], and [truth] (p. 1114). I do not know if this distinction exhaustively addresses the relevant differences, but it does seem to be a sound approach.

While it is not impossible that the mental representation of concepts may include sensorimotor information, it does seem impossible that sensorimotor information alone can constitute certain concepts. I should point out that (while I am not ruling out perceptual concepts; see the end of this section) the difficulty of how sensorimotor information alone can constitute concepts persists for concrete concepts as well: knowledge of what a hammer looks like and how to use a hammer and how using a hammer or being hit by a hammer feels does not constitute the concept of a hammer, and this is because hammers as concepts are not purely sensorimotor objects. Pecher (2018) and Pecher and Zeelenberg (2018) review experimental evidence showing that sensorimotor processing of concrete objects is task-dependent, and activation of sensorimotor representations may not be automatic. So it should not be assumed automatically that even concrete concepts are built wholly from perceptual representations.

The other major objection relates to complications in the interpretation of the aforementioned experimental data: the many empirical results that have been offered as evidence of a perceptual representational format, while consistent with a modal hypothesis, are not necessarily inconsistent with an amodal hypothesis. There are two main aspects to this.

Firstly, it is not clear whether the sensorimotor cognition seen empirically in association with conceptual processing is constitutive of the mental representation of concepts or is secondary (i.e. simply activated in association with the mental representation of those concepts and perhaps involved in processing them) (Machery 2007; Hauk and Tschenstcher 2013; Leshinskaya and Caramazza 2014; Mahon 2015a, 2015b; Mahon and Hickok 2016).

Associative activation is a very general brain mechanism, and so it would be expected that the mental representation of a concept would trigger other—distinct—activity in sympathy with it. The occurrence of sensorimotor activation in contexts where one would expect a concept to be mentally represented is not proof that that activation *is* the mental representation. Processing of concepts is not the same thing as representation of concepts.

It has been pointed out (e.g. by Hauk and Tschentcher (2013), Leshinskaya and Caramazza (2014), Mahon (2015a, 2015b), and Mahon and Hickok (2016)) that associative activation has not been ruled out as an explanation of the neuroimaging data. Activation of sensorimotor areas of the brain is seen on the order of 200 milliseconds after the presentation of the stimulus (Kiefer and Pulvermüller 2012; Hauk and Tschentcher 2013; Galetzka 2017), which has been argued (e.g. Kiefer and Pulvermüller 2012) to be too fast for it to represent secondary imagery inspired by and occurring after conceptual representation. However, it is not at all too fast to be resultant from associative spreading of activation: for example, accessing the phonological form of a particular word in speech production (thought to be effected by activation spreading) is estimated to take something on the order of 200 milliseconds (Indefrey and Levelt 2004; c.f. Pickering and Garrod 2013); in speech comprehension, identifying a phonological form from an acoustic signal is estimated to take something on the order of 100 milliseconds (Friederici 2002; c.f. Thompson and Kielar 2014).

It is a further reach to draw conclusions about the format of mental representations from the brain regions they are found in: specification of a brain region cannot be taken as specification of format; while neuroscience has been able to associate particular brain regions with certain (generally vague) functions, nothing is known directly about the format of mental representations and how it may or may not correlate with brain region (c.f. Leshinskaya and Caramazza 2014: 126-130; Mahon 2015a, 2015b: 174-175).

Secondly, even if mental representations of concepts are built from sensorimotor representations, the sensorimotor representations may be so indirectly involved that in terms of format, the mental representation of concepts is not perceptual in any meaningful sense. The complicating theoretical factor is that as soon as "perceptual" is allowed (as it must be—otherwise, constituting concepts is totally hopeless) to go beyond purely sensory representations and include various levels of processing of sensory information, it becomes difficult to disentangle representations that should properly be considered sensorimotor or perceptual from more abstract representations that should not.

At low levels, there are representations of purely sensory information, and at higher levels, there are various levels of processing (i.e. interpreting it). For example, spatial representations can be built from visual information or auditory information or tactile information. At these higher processed levels of representation, modality-specificity is lost, and at the same time abstraction is gained. The problem here is: At what point is it still proper to call the representation perceptual? Probably, the level(s) of interpretation responsible for effecting visual illusions, for instance, can safely be considered to be perceptual; and that involved in recognizing that something can be either inside or outside a container, but not both, can safely be considered to be non-

perceptual[57]; the case of spatial representations is less clear. It does not seem possible for pure sensory information alone to instantiate concepts; at the very least, it would have to be interpreted. So this question is particularly important for conceptual representations.

Concerning higher levels of processing of perceptual information, one proposal (e.g. Barsalou 1999) is that at higher levels of abstraction, various pieces of sensorimotor information would be integrated in *convergence zones* (hypothesized by Damasio (1989)) (perhaps in the inferior parietal lobe; c.f. Galetzka 2017; Montefinese 2019). The representations constructed in these convergence zones are generally characterized as multimodal, as opposed to amodal[58]. The distinction between multimodality and amodality itself is a valid one: for instance, tactile information and visual information (about a ball, for example) can be integrated, and the resulting representation be considered multimodal because the content represented is sensory[59]. But if, on this definition, multimodality is distinguished from amodality, the problem of interpretation is not solved by integration, since there is no abstraction. For example, a spatial representation goes beyond sensory content (it includes abstract interpretation), and therefore is not cleanly separated in content from amodality. It appears that any version of multimodality that includes abstraction or interpretation becomes hard to

[57] This example is actually from Lakoff (1987: 272 (also see pp. 365-367); c.f. Hauk and Tschentscher 2013: 2), who uses it to argue a somewhat opposite point: that understanding of logic is rooted in perception. It is certainly possible that perceptual experience with objects being either inside or else outside containers contributes to the acquisition of an abstract concept of mutual exclusivity, but when it comes to the format of mental representation, it seems absurd to posit that it is perceptually represented, unless one pre-assumes that all representations are perceptual. For a view of image schemas as non-perceptual abstractions, see Pecher (2018).

[58] Barsalou (1999: 637) does concede that it is possible in principle for these representations to be amodal.

[59] The format of this content is important; I am not referring to content in simply a 'thematic' sense here. As Mahon (2015a) remarks, "I may represent the knowledge that fires are red in a format that has nothing to do with the sensory systems that process redness. Thus, the content of my representation that fires are red is *about* visual information—but that doesn't make my representation that fires are red at all 'embodied.'" (p. 11; also see Leshinskaya and Caramazza 2014: 124-125).

distinguish from perceptually-grounded amodality. (See also Markman and Stilwell 2004: 399-400; Haimovici 2018: 2-4.)

Another, similar, proposal (Rogers et al. 2004; Patterson et al. 2007; Lambon Ralph et al. 2009; Lambon Ralph 2014; see also Dove 2016; Galetzka 2017; Montefinese 2019), known as the *hub-and-spoke model*, hypothesizes that information from various modalities (including perceptual information) is integrated in a central 'hub' (in the anterior temporal lobes), and used to build the mental representation; the format of this representation, however, is amodal. This can be considered a weak form of the modal hypothesis, which, as Mahon (2015a) argues, becomes very much like an amodal hypothesis. The distinction, then, is between whether (for the weak modal hypothesis) sensorimotor processing forms (at least part of) the basis for more abstract representations (which are themselves non-perceptual) or (for the amodal hypothesis) the sensorimotor processing is activated in association with the abstract representations, without (being involved in) constituting them.

At any rate, certain conceptualizations clearly involve abstraction beyond what sensorimotor systems are capable of. A mathematician thinking about a 10-dimentional cube, for instance, cannot visualize it in an ordinary sense, and in any sense in which they can visualize it, either the "visualization" would involve a generalization of sensorimotor processing in which it is abstracted to beyond a purely perceptual function, or there is a separate system that is not purely perceptual that is responsible. So it does seem that non-perceptual processing must be involved at some point, for at least some concepts.

A point of attractiveness for the modal hypothesis is that it provides a potential answer to the problem of how mental representations can carry meaning, known as the *grounding problem* (from Harnad 1990). The answer it gives, of course, is that mental representations do so by depicting perceptual information (c.f. Harnad 1990; Barsalou 1999; Kaschak et al. 2014; Galetzka 2017; etc.); how mental representations of perceptual images come to be meaning-bearing is not so mysterious[60]. Adherents of the modal hypothesis often assume that the only way mental representations can represent

[60] How they are neurally instantiated is another matter, of course.

meaning is through depicting sensorimotor content, and so mental representations must be grounded perceptually. For example, Kaschak et al. state "if linguistic symbols are not grounded in systems of perception and action planning, it is difficult to understand how they could convey meaning" (2014: 120). Part of the reason for this way of thinking may be a conception of amodal representations as abstract, arbitrary symbols; certainly, a system consisting of abstract symbols and nothing else cannot carry meaning. But amodal representations, if they are to be conceived of as symbols, do not have to be *empty* symbols (c.f. Mahon and Hickok 2016: 945-946; Haimovici 2018: 5-6). So the question of how amodal representations are imbued with meaning is just a mystery and not a fatal problem. Additionally, since (as discussed above) even perceptual representations have higher levels of interpretation and abstraction, which seems to require non-perceptual information, modal theories (at least formulations that are not unviably strong) are not entirely free from the grounding problem (c.f. Mahon 2015a: 8-10; Pecher 2018).

So, the question of what the format of the mental representation of concepts is is very much an unresolved issue. I have discussed the modal thesis of mental representation with cautious criticism, remarking on various fundamental challenges it faces. I do not reject it absolutely, since I cannot rule out that perceptual representations contribute to conceptual representations. My stance is agnostic between amodality and weak modality, and on some points I lean more toward amodality (for example, I am reluctant to assume that concrete concepts are represented perceptually) and on others (like the following) toward weak modality. I should point out that an amodal view does not entail the denial of the use of imagery or simulation in cognition (c.f. Machery 2007). It is also likely that emotional representations are highly active in cognition in general (c.f. Vigliocco et al. 2014). Additionally, it seems that there are certain cognitive objects that can only be considered concepts whose content is entirely perceptual: for example, in order to perceptually categorize something, there must be some concept against which it is categorized, and this concept cannot be amodal—it must be perceptual; one can have two perceptual images and compare them, and in order to refer to one later, it must be a concept. I call these *perceptual concepts*; they are (a small) part of the topic discussed in the next section.

2.2.3 Conceptual lexicalization

The goal of this section is to introduce the notion of *conceptual lexicalization*. In linguistics, *lexicalization* refers to the phenomenon of having a particular concept encoded in a lexical item. There is a great variety across languages in what is lexicalized and what is not, such that speakers of one language may be surprised by the exoticness of what sorts of concepts other languages have or do not have words for. This is the subject of the first half of this section. The point of this is: essentially anything conceivable can be turned into a concept and referred to through a label. The ability or process by which this is done is what I call *conceptual lexicalization*; conceptual lexicalization is discussed in the second half of this section. Conceptual lexicalization is important for the organizing of language-independent thought for expression in language.

2.2.3.1 Variety in lexicalization

Languages lexicalize weird concepts. They are weird in the sense that each concept is idiosyncratic: it is itself and not any other out of the whole space of possible concepts. A concept being weird is of course relative: speakers of language A may wonder why speakers of language B would ever come up with a word for such a weird concept, but speakers of language B may wonder why language A does *not* have that concept. So the presence of a concept in one language demonstrates at the same time its lack in others.

Evans and Levinson (2009: 435) remark: "Languages may lack words or constructions corresponding to the logical connectives "if" (Guugu Yimithirr) or "or" (Tzeltal), or "blue" or "green" or "hand" or "leg" (Yélî Dnye [a Papuan language]). There are languages without tense, without aspect, without numerals, or without third-person pronouns (or even without pronouns at all, in the case of most sign languages). Some languages have thousands of verbs; others only have thirty (Schultze-Berndt 2000)."

Differences can be subtler, but nonetheless significant. While English *float* can express movement (as in *The bottle floated into the cave*), the Spanish near-equivalent

flotar does not express movement; thus, in Spanish, one would have to say something like 'The bottle went into the cave floating' (*La botella entró a la cueva flotando*) (Talmy 2000: 49). Likewise, while one may say *I kicked the door open* in English, in French, one must say something like 'I opened the door with a kick' (*J'ai ouvert la porte d'un coup de pied*) (Levin and Rappaport Hovav 2019: 396). Similarly, the meaning of *off* in English enables one to say things like *He filed the serial number off*, but there is no analogous structure in French; instead, it is expressed as something like 'He removed the serial number with a file' (*Il a enlevé à la lime le numéro de séri*) (Levin and Rappaport Hovav 2019: 403, citing Green 1973: 273).

Chinese *cái* and *jiù* express that an event is remarkably late or early, respectively. They are not entirely dissimilar to *finally* and *already* in English: one may say *I finally finished it*, implying that it took a long time to do so, or *You finished already?*, expressing surprise that it has been done so soon. But if one wishes to express that one finished it in an hour, and this was remarkably fast, one cannot say *I already finished it in an hour*. However, this can be done with *jiù*: *wǒ huā yī xiǎoshí jiù wánchéng le* (lit. 'I take one hour *jiù* finish'). [61] The same goes for the case when the hour was remarkably slow: *wǒ huā yī xiǎoshí cái wánchéng* (lit. 'I take one hour *cái* finish') expresses it, but *I finally finished it in an hour* does not work.

The names for objects that correspond to English *cup, glass*, *bottle*, *jar*, and *container* are divided up differently in languages such as Hebrew, Japanese, Russian, Spanish, and Chinese (Malt and Masjid 2013: 587-588).

Languages often make distinctions that seem bizarrely specific and obscure from the point of view of other languages.

The Alaskan language Yup'ik has a complex demonstrative system, that instead of distinguishing just between *this* and *that*, distinguishes between referents that are in more-accessible locations and less-accessible locations, and further between referents that are close to the speaker, higher than the speaker, lower than the speaker, at the same level as the speaker and away from them, and inside or outside from the speaker (here, more accessible is inside the speaker, and less accessible is outside the speaker)

[61] It should be noted that Chinese has other expressions (*yǐjīng* and *dōu*) that more directly correspond to English *already*.

(Tamura 2014, 2017). Tzeltal has terms that indicate ‘of flat items, arranged in vertical stack’ (*latz’al*) and ‘be located in bulging bag’ (*chapel*) (Evans and Levinson 2009: 435); the Californian language Karuk has a term that indicates ‘in through a tubular space’ (Evans and Levinson 2009, citing Mithun 1999: 142). The Californian language Atsugewi has verb roots with meanings like ‘for a small shiny spherical object (like a round candy) to move/be located’ (-*lup*-) and ‘for a slimy lumpish object (like a toad) to move/be located’ (-*caq*-) (Talmy 2000b: 58).

Whereas English has only *carry* and *hold*, Chinese has a large number of verbs that distinguish between different manners of carrying or holding, which—unlike English—do not distinguish between whether the object is being carried or simply held (Saji et al. 2008; Saji et al. 2011). Examples are shown in Table 3 below.

Table 3: Carrying/holding verbs in Chinese (adapted from Saji et al. 2011: 50)

Verb	Manner	Example of typical object
bào	carry/hold an object in both arms	(stuffed animal)
bèi	carry/hold an object on one's back	(rucksack)
dǐng	carry/hold an object on the top of one's head	(bowl)
duān	carry/hold an object by hand, keeping the object horizontal	(bowl with water)
jiā	carry/hold an object under one arm	(square bag)
jǔ	carry/hold an object over one's head	(square box)
káng	carry/hold an object by resting it on one's shoulder	(pipe)
kuà	carry/hold an object by hanging it on one's shoulder	(tote bag)
līn	carry/hold an object by dangling it with one hand	(plastic bag)
ná	carry/hold an object with one hand	(plastic bottle)
pěng	carry/hold an object cautiously in both hands	(bouquet)
tí	carry/hold an object by dangling it around one's arm	(handbag)
tuō	carry/hold an object with one's palm(s)	(tray)

Some languages have pronouns that distinguish between even- and odd-numbered generational distance, or express relationships between the speaker, hearer, and referent simultaneously, like 'the one who is my mother and your daughter, you being my maternal grandmother' (Evans and Levinson 2009: 436, citing Evans 2003)

Strikingly exotic concepts can be encoded in lexical items. A canonical example from English is s*chadenfreude* ('pleasure at another's suffering')[62]. Evans and Levinson (2009: 435, citing Osada 1992) give the examples of the Mundari (an Austroasiatic language spoken in eastern India) onomatopoeic terms *ribuy-tibuy*, meaning 'the sound,

[62] *Schadenfreude* is analyzable in German (<harm> + <joy>), but unanalyzable in English.

sight, or motion of a fat person's buttocks rubbing together as they walk', and *rawa-dawa*, meaning 'the sensation of suddenly realizing you can do something reprehensible, and no-one is there to witness it'.

The variety glimpsed through these examples demonstrates two things: that the content of concepts is not very restricted, and that essentially any concept can be lexicalized. This leads us to the second part of this discussion.

2.2.3.2 Indexed concepts and conceptual lexicalization

We will now discuss the cognitive ability/process I call *conceptual lexicalization*, which is key for the formation of linguistic representations and organizing language-independent thought for expression in language.

To start with, it is readily apparent that any arbitrary concept can be taken and turned into a word [63] —i.e. lexicalized; some examples were mentioned above. Underlying this fact is the cognitive ability to perceive some bit of conceptual material as a unit and index it so it can be referred to as a concept (or 'conceptual unit')[64]. There are two components to this: the unification of conceptual material[65] and indexing of that unified material. The point of the index is to enable the retrieval of a specific chunk of conceptual material. The index uniquely identifies the concept.

One might naively think that concepts are (or can be) accessed through their linguistic names, and therefore the index *is* the overt (phonological or orthographic)

[63] This extends to linguistic elements in general, though the kinds of meanings constructions can express are rather restricted (c.f. Section 2.2.1.1 (The (non-)distinction between lexicon and grammar) above).

[64] There is no settled or agreed-upon general characterization of concepts, either concerning their function or their mental representation. That matter is outside the scope of this dissertation, and I will not attempt to resolve it here. For recent general discussion of concepts, see Machery (2009, 2010), Carey (2009, 2011), Margolis and Laurence (eds.) (2015), and Hampton and Winter (eds.) (2017).

[65] I will not have more to say about this because I do not know what cognitive mechanisms are involved, but we do know that such a process has to occur, because even if the content of a concept is formed from distinct bits, it still functions as a unit.

lexical form; but these are quite distinct. For instance, in the tip-of-the-tongue phenomenon, language users know precisely what word they want, but cannot access its externalizable form. That is to say that they have identified the particular meaning they want—i.e. the unique, indexed concept—but are (temporarily) clueless about the externalizable form of the associated word (this will be discussed in more detail in the next section (The architecture of language)). Secondly, externalizable forms do not uniquely identify the concept: for example, in homophony (e.g. *I like cookies* vs. *This cookie is like that cookie*) or polysemy (e.g. *My kid depends on me* vs. *That depends on the situation*), a single externalizable form is associated with several distinct concepts. Also, as will be discussed immediately below, externalizable forms are not at all required for indexing.

Indexed concepts do not have to be attached to a word or some other overt form. Many indexed concepts do not have linguistic 'names'—i.e. there is no word etc. for them. An easy-to-understand example, offered by Akira Machida (personal communication), is the phenomenon of one person meeting someone in a hallway and stepping aside to let them pass; but they do the same and both are now again in each other's way. This does not have a word, but can be mentally referred to at will: it is an indexed concept. Perceptual concepts are generally of this type; particular perceptual images[66] can be mentally referred to despite not being associated with any linguistic element. Nuyts (2012: 330) mentions the taste of garlic, jalapeños, and cilantro, which can be distinctly conceptualized and recalled but have no linguistic labels (and cannot really be described satisfactorily). The same is true for the particular feeling of being not-quite awake after waking up too early, the particular feeling or emotional reaction elicited by seeing a particular shade of blue sky, etc. One can identify a person as someone one has seen before without knowing their name (or giving them a made-up one, like *funny glasses guy*), and one can have in mind a particular person while failing to recall their name.

Now, indexed concepts *could* be associated with words. Words are simply indexed concepts with a linguistic symbol attached to them. Where *lexicalization* is the

[66] The image may be from a unique experience or be abstracted over several instances, but this does not matter for the present point.

assignment of a linguistic label to a(n indexed) concept such that it can be referred to in language, *conceptual lexicalization* is the assignment of an index to a concept such that it can be referred to cognitively. This is why I call it conceptual lexicalization.

Language depends very fundamentally on conceptual lexicalization. The ability of conceptual lexicalization underlies the system of linguistic symbols that refer to particular concepts. Conceptual lexicalization establishes retrievable concepts, which can then be linguistically encoded. Also, indexing of concepts is essential for language use, since it enables the access of linguistic elements; this is discussed more in the following section.

2.3 The architecture of language

Our next subject is what is known as the architecture of language: the theory of how the system of language is organized and structured—what its components are and how they are connected. In its most basic sense, language as a system consists in a pairing of symbol and meaning. There are a variety of theories, which differ on the details of how this is effected. Our understanding of the architecture informs how we approach our investigation of the expression of thought in language, as will be expounded on in the next section.

First, a clarification on the terminology of many of the theories reviewed here is needed. The pairing of symbol and meaning is very often (but by no means exclusively) talked about—by researchers across linguistics, from generativists like Chomsky (2007 etc.) and Jackendoff (2011 etc.) to cognitivists like Langacker (2008 etc.) and Evans (2009)—as pairing of "sound" or "phonology" and meaning. Now, what is meant here by "sound" or "phonology" is not actually sound or phonology. Firstly, to the extent that sound is involved, what is relevant is of course the mental representation of phonological form rather than the physical sound itself (c.f., e.g., Langacker 2008: 15 note 10). Secondly, and perhaps more substantially, language is not restricted to the modality of speech: it can be signed or written, so the overt form is not necessarily phonological. So "sound/phonology" refers to the overt form rather than specifically the

phonological form. Researchers are sometimes explicit about this. For instance, Langacker (2008: 15) clarifies: "Under the rubric *phonological structure*, I include not only sounds but also gestures and orthographic representations. Their essential feature is that of being overtly manifested, hence able to fulfill a symbolizing role." So "sound" or "phonology" essentially refers to overt form in general.[67] Also, some theories recognize that linguistic symbols include constructions, which do not correspond directly to overt forms (e.g. Croft 2001; Evans 2009), so (depending on the theory) the linguistic symbol paired with some meaning is not necessarily overt.

Chomsky and colleagues (e.g. Chomsky 2007; Berwick and Chomsky 2011) conceive of language as essentially consisting of a collection of words (the lexicon) and procedures for computing conceptual structure; the lexicon forms the "atoms of computation". The current Chomskyan paradigm, the Minimalist Program, reduces syntactic computation to a single operation, Merge, which has two variants. External Merge effects syntactic embedding by taking two elements and binding them into a set (a, b ⇒ {a, b}). Internal Merge effects syntactic displacement through copying-and-deletion by binding a set with one of its constituents (a, {a, b} ⇒ {a, {~~a~~, b}}). Chomsky and colleagues emphasize that Merge can take as its input the output of a previous iteration of Merge, which can be repeated ad infinitum; this reflects the fact language does not impose limits like a cap on the number of syntactic embeddings. A linguistic expression is described as a syntactic object that is mapped to a phonological form and to a semantic interpretation (Hauser et al. 2002; Chomsky 2007; Berwick and Chomsky 2011; etc.). This architecture is shown below in Figure 1.

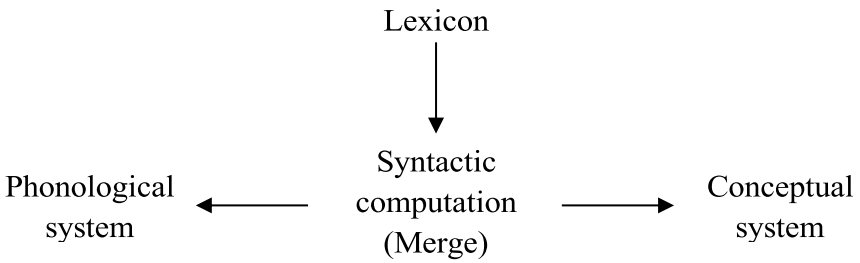

Figure 1: The Minimalist architecture

[67] One wonders, then, why a different term is not used (one being me).

For our purposes, the principal problem with this model is that it gives a strict separation between the lexicon and syntax. We are concerned with language as a system for expressing meaning; the point of such a language system is not to figure out what syntactic positions elements go in—rather, syntactic structures are part of the "atoms" with which meaning is expressed. The unsuitability of this separation has been pointed out by Jackendoff, Langacker, and others (as discussed in Section 2.2.1.1 (The (non-) distinction between lexicon and grammar)), who have articulated their own visions of the architecture of language.

Jackendoff (1997, 2002, 2007, 2011, 2017) has criticized the Chomskyan view for describing linguistic compositionality as purely in the syntax[68]. He argues that such a theory cannot account for phenomena like the fact that *Jill slept until the alarm went off* indicates a single act of sleeping, but *Jill jumped until the alarm went off* indicates repeated jumping, and *Jill jumped when the alarm went off* indicates a single jump, where differences in meaning do not correlate with differences in syntax (2007: 45). Another criticism (discussed earlier) is that constructions also contribute to meaning: the 'lexicon + syntax' model is inadequate.[69] He proposes instead what he calls the parallel architecture, which views syntax, semantics, and phonology as independent combinatorial systems that are connected with "interface rules". This is shown below in Figure 2. Though it is partially similar to the Chomskyan architecture in that the description is in terms of separate systems of syntax, semantics, and phonology, in the parallel architecture, combinatoriality in semantics and phonology do not derive from combinatoriality in syntax. The interface rules are the lexical items (in the broad sense described above in Section 2.2.1.2) of the language. Words, morphemes, and constructions can be thought of as rules that link bits of conceptual structure, syntactic structure, and phonological structure. Jackendoff notes that linguistic elements do not necessarily contain all three: in his analysis, *hello* and *yes* have phonology and semantics, but no syntax; *it* in *It's noisy in here*, *do* in *I didn't see her*, and *of* in *a*

[68] He refers to this conception as "syntactocentrism" (2002, 2007).

[69] Jackendoff (2011) also criticizes the conception of Merge as the basic combinatorial principle, but for Jackendoff, this is not really an architectural issue, so I will not discuss it here.

picture of Bill have phonology and syntax, but no semantics[70], constructions like that in *Joe knitted the morning away* have syntax and semantics, but no (or only partial) phonology, and nonsense phrases like *bibbity-bobbity-boo* used for metrical purposes in music have phonology only, and no syntax or semantics (2011: 609-610, 2017: 193).

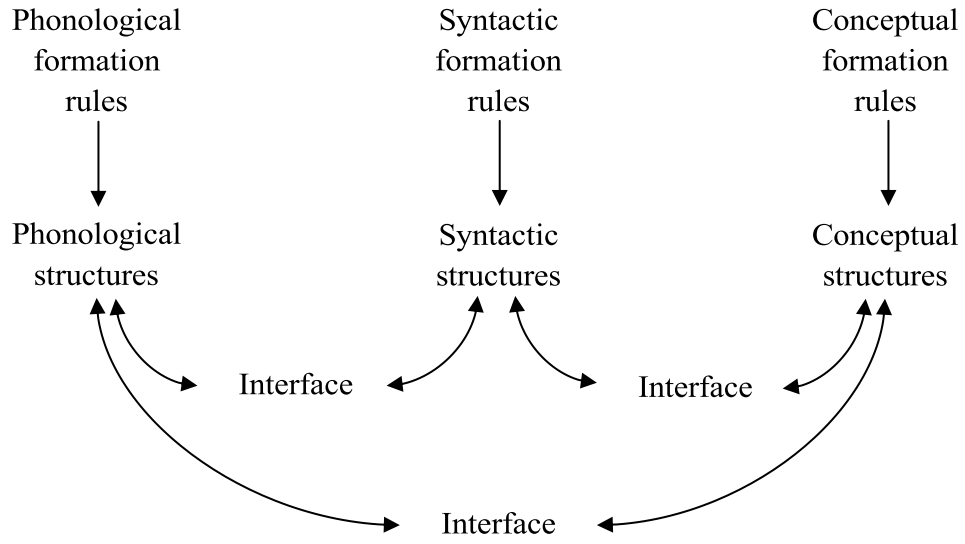

Figure 2: Jackendoff's parallel architecture (2011: 609 etc.)

While I agree with Jackendoff that meaning in linguistic expressions should not be described as being derived from syntactic structure, I do not think that meaning is separate from (what can be described as) syntactic structure. This is because everything that can be described as syntax—constructions, tense, etc.—is meaningful. It is not possible to manipulate syntactic structure without at the same time manipulating semantic structure.[71]

The notion that grammar is inherently meaningful is central to Langacker's approach, which is representative of cognitive linguistics in general. Langacker (1987, 2008, 2009b; etc.) describes linguistic expressions as assemblies of symbolic structures

[70] I do not think this analysis is at all tenable for *of*, but *it* and *do* are less clear; see footnote 72 below.

[71] This disjoint may lie in Jackendoff's apparent conflation (discussed earlier) of the thought to be expressed (which *is* generally separate from any syntactic structure) and the meaning actually encoded in the expression (which is not).

(Figure 3 below). A symbolic structure (denoted as Σ in the figure) is a pairing of two components: a *semantic pole* (S; i.e. the meaning) and a *phonological pole* (P; i.e. the overt form). Symbolic structures can be words, morphemes, and constructions; there is no principled distinction between these. Thus, "lexicon, morphology, and syntax form a continuum fully reducible to assemblies of symbolic structures" (2008: 15). Symbolic structures can be combined into larger structures; these constitute assemblies. Single-element expressions can be considered a special case of an assembly (Langacker uses the term *degenerate symbolic assembly*). Significantly, Langacker does not recognize a separate syntax; there are likewise no descriptions of interfaces between separate systems. The generalization in the notion of "assemblies of symbolic structures" subsumes syntax as part of a single system whose purpose is the expression of meaning through form.

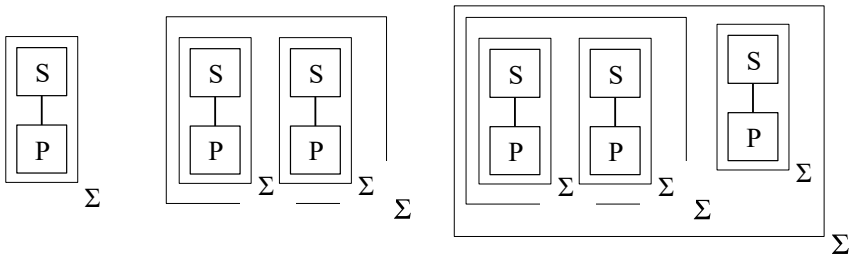

Figure 3: Langacker's architecture (2008: 15)

This conception is closer to the characterization we want, but the description of the notion of symbolic structure needs more precision. The relationship between the semantic poles and phonological poles is not symmetrical. Some Ps are empty (as noted by Jackendoff): some linguistic elements do not correspond to any particular overt form (e.g. the construction in *I {gave/mailed/threw/etc.} him the ball* is not—unlike, for

instance, *ball*—denoted by an overt symbol). But no S can possibly be empty[72]. Meaning is primary, particularly for language production. Evans (2009, particularly pp. 97-98) partially solves this problem, by differentiating between what he calls *phonetically overt vehicles* (i.e. overt linguistic symbols) and *phonetically implicit vehicles* (i.e. non-overt linguistic symbols). But there is another part to this problem, which I alluded to earlier in the discussion on indexing of concepts in Section 2.2.3: the semantic content of a linguistic element is uniquely identified at a level prior to the involvement of any linguistic symbols.

Levelt et al. (1999) almost solve this second point. Their model (which is specifically on lexical selection, so it focuses on words) posits three levels of representation: the *conceptual stratum*, the *lemma stratum*, and the *form stratum* (see

[72] I state this with the caveat that analyzability can make the situation a bit messy. Earlier, I gave the example of a language user analyzing *metaphor* as being made of the morphemes <meta> and <phor>, but not knowing what they mean. In this case, one can say that these low-level morphemes have Ps, but S does not inhere until the higher level where they are joined (cases like *look up* are similar, except 'residues' of other meanings of the components *look* and *up* (ones that are viable independently) may contribute to the higher-level meaning). But in actual language production, *metaphor* would have to be accessed as a whole in this case—and it does have an S.

Jackendoff describes *it* in *It's noisy in here* has having no semantics (2017: 193); Langacker describes it as meaningful (denoting an abstract setting) (2008: 390, 451-453). This may also relate to analyzability: some language users may analyze *it* as expressing an abstract setting (i.e. Langacker's analysis); others may not attribute any particular meaning to *it* (i.e. Jackendoff's analysis), but necessarily do attribute meaning to the larger construction it is a part of.

Also, the other case that Jackendoff identifies as having no meaning should be commented on. Jackendoff says nonsense phrases like *bibbity-bobbity-boo* (ibid.) that are used to fill up metrical space in songs have phonology, but no semantics (or syntax). I would say they are in fact not empty semantically, if they are actual linguistic elements. In general, 'sound effects' can be lexicalized (c.f. conceptual lexicalization, discussed above), as in *He went waaa*; here, *waaa* can be intended as onomatopoeia for the sound that he produced, or it can be some more abstract iconic 'depiction' of whatever action he performed. In these cases, the S is certainly not empty. The specific case of filler in song lyrics that he refers to I would classify as non-linguistic: the forms in question certainly have linguistic phonological structure, but they function as musical elements rather than linguistic ones.

Figure 4). The first and last of these are quite straightforward: the conceptual stratum contains the concept associated with a word, and the form stratum contains the form of the word. The lemma is a strange amalgam. Part of the function of this level of representation is to serve as a marker for the conceptual content, such that it can be matched to a particular linguistic element. Levelt and colleagues discuss the tip-of-the-tongue phenomenon as suggesting the existence of this level of representation: experiencers of the phenomenon are not suffering from an inability to find a suitable expression; they know they want to say a particular word, but cannot access its phonological form. One would expect from this fact a level of representation in which the meaning is not just coincident to that of the linguistic element in question, but specifically identified as belonging to it. If a particular linguistic element has been selected or accessed, the mental representation involved will contain the conceptual content of its meaning and a marker identifying that content as uniquely belonging to that particular element. So this function is identical to that of what I have termed the *index*. But another role of the lemma is to contain the syntactic information for the word; here, morphological inflections are treated as "diacritic parameters". So, for example, in *escorting*, the *–ing* inflection is described as a diacritic attached to the lemma *escort*. They describe the semantics of a word as represented at the conceptual level and the syntax of the word as represented at the lemma level; but such a segregation is impossible, since (as they acknowledge) any information related to grammatical encoding must first be determined conceptually. It is unclear to me why "semantics" and "syntax" is split between levels and why syntactic information (alone) should be placed together with the index. As I have discussed, I do not think a separation between semantic information and syntactic information is well-motivated.

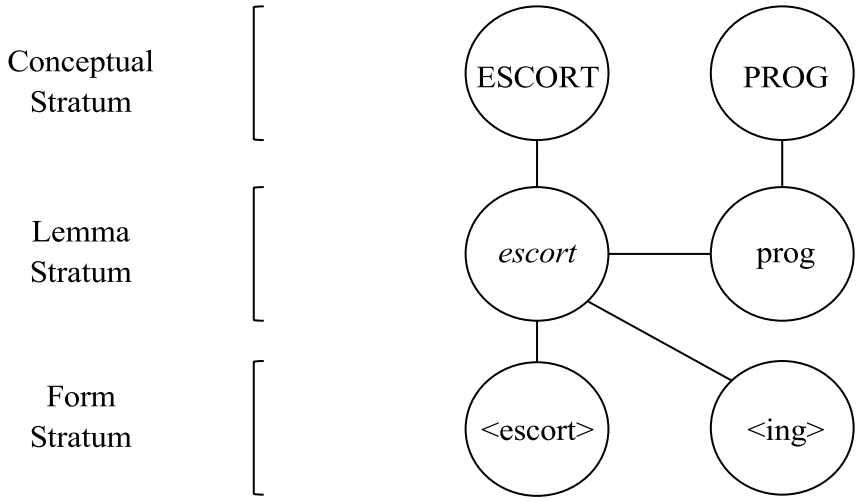

Figure 4: Levelt and colleagues' architecture (adapted from Levelt et al. 1999: 4, 12)

I will now propose my own architecture. My conception is shaped by the issues discussed above. The basic notion is that a linguistic element is an indexed concept with a linguistic symbol attached to it. This is shown in Figure 5. There is no principled distinction between syntactic information and semantic information here: it is all meaning. While there is definitively a single level at which the meanings associated with linguistic elements are represented (i.e. no "semantic level" and "syntactic level"), I do not know if the index is better described as a separate level of representation (like Levelt and colleagues' lemma) or as a component of the representation of the concept. The linguistic symbol can be overt (e.g. a word) or non-overt (e.g. a construction whose linguistic symbol is a syntactic pattern). It should be noted that the externalizable form (phonology etc.) is not central to the representation of the linguistic element: what it means to access an element is to access the indexed concept. When a language user is in a tip-of-the-tongue state, they have accessed the linguistic element in question, but not

the externalizable form. So mentally accessing the element does not require access of the symbol; the indexed concept is the crucial part of the element[73].

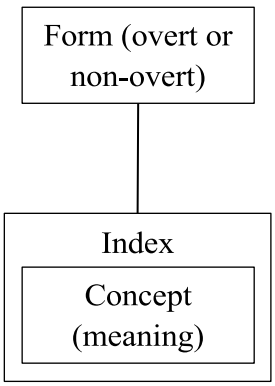

Figure 5: Representational architecture of a linguistic element

Linguistic representations are arrangements of elements. A single element is a special case of an arrangement, so this applies to expressions or local bits of expressions that consist of single elements as well[74]. The meaning of the expression is constituted by the arrangement of component meanings; that is to say that the arrangement of component meanings yields some effective meaning, which is the meaning of the expression. However, the overall meaning is not simply a structured list of component meanings. Meaning construction and the ways that bits of meanings in an arrangement interact with each other are complex; this is a major topic in the next two chapters.

The architecture we assume has bearing on how the outline of the process of expressing thought in language is characterized. That is the subject of the next section.

[73] I do not mean to claim that the level of representation of the externalizable form is absolutely subservient to that of the meaning. Phonological information can influence accessing of linguistic elements, through spreading of activation (as seen in, for instance, phonological priming) (see Dell et al. 2014), so there can be a slight degree of interplay between form and meaning. See Section 4.2.2 (Accessing elements) for more discussion.

[74] An argument can be made that single-word expressions involve a construction as well (e.g. *Food!* said to bring someone's attention to food can be analyzed as having a construction in addition to the word; so can *Hello*), but whether or not true single-element expressions exist does not really matter for the theory I present here.

2.4 Schema of the process of expressing thought in language

Here, we will be outlining the schema of the process of expressing thought in language: where it starts and where it ends and what basic steps it contains. The material discussed in previous sections motivates the approach laid out here. The discussion in this section sets up the problem that the theory laid out in the succeeding chapters has to account for. Essentially, the process starts with one representation of meaning (the original thought) and ends with another—different—representation of meaning (the linguistic expression). The linguistic expression is an approximation of the original thought, and the goal of the cognitive process we are concerned with is to derive it.

Before going into the specifics of the schema of the process, there are some preliminary points that should be noted.

Actual on-line production of language is often 'messy'. The theoretical ideal would be that the linguistic expression is the best possible approximation of the thought, but this is often not the case in reality. As noted above in Section 2.1.5 (also see Section 4.2.2), more mentally-accessible (habitually or commonly-used or otherwise primed) linguistic structures are often favored, so the produced expression is not necessarily the closest possible fit to the original thought. Essentially, the linguistic expression is a 'good enough' approximation of the original thought. Presumably, different people have different tolerances regarding how faithful the expression must be to be acceptable, which depend also on the situation of language use (i.e. in some cases language users may be more careful about their wording than in others). So the end product is a 'suitable' linguistic expression, where 'suitable' is whatever is accepted by the language user. This will generally be quite close to the original thought, but not necessarily the best possible approximation of it.

Also, the thought to be expressed often does not emerge fully formed as a unit ready to be turned into a sentence: often, a language user will, for example, start speaking with only part of their thought formed, and develop the thought as they speak; they may change their thought halfway through as they work out their idea; and they may change what is to be expressed for pragmatic reasons (c.f. Bock and Ferreira 2014:

30-32; Hartsuiker 2014, particularly pp. 427-428; also see Section 5.3). Language production is *incremental* to varying degrees (see Konopka and Brown-Schmidt 2014: 6-8 for discussion). Regarding variability in the size of the chunk of conceptual material that gets expressed at once, Konopka and Brown-Schmidt (ibid.: 16) suggest that more complex thoughts may require more holistic processing.

Additionally, some remarks on the delimitation of the scope of our problem need to be made. There is a consensus—one that I agree with—among researchers on speech production (e.g. Levelt 1989, Levelt et al. 1999; Bock and Griffin 2000; Pickering and Garrod 2013; Konopka and Brown-Schmidt 2014; Bock and Ferreira 2014; Dell et al. 2014) that the overall process of speech production consists of three basic stages: it begins with coming up with an idea to be expressed, progresses through the formulating of the linguistic expression, and ends with the physical articulation of that expression. The problem I am considering in this dissertation is the purely mentally-internal process of coming up with a linguistic expression of a thought. The expression does not need to be externalized in a physical utterance (though it can be, of course): it can be used in thinking to oneself in an 'inner voice', or thinking of something that one may potentially say or write; externalization is a separate problem. The stages subsequent to the formation of the expression (as an object that represents meaning) are outside the scope of our problem, since they apply after the expression has been set; the thought has already been expressed. Articulation itself, of course, is wholly outside our scope. Phonological/phonetic and orthographical encoding are also not part of our problem, since they are encodings of an already-decided linguistic expression. If one intends to speak, one needs to mentally represent the phonological form of the expression. But constructing a mental representation of the phonology is not integral to the creation of the mental representation of the linguistic expression in the sense we are concerned with. This is well-illustrated with the tip-of-the-tongue phenomenon, since there the expression is decided without any representation of phonology. So while phonological encoding is often done automatically—i.e. done even when (external) speech is not the goal, e.g. as in inner speech (c.f. Alderson-Day and Fernyhough 2015; Langland-Hassan 2018; etc.)—, phonology is unimportant (except for priming effects, as noted above). Additionally, the formulation of the thought is not part of our problem.

In communicative situations, pragmatic concerns (e.g. understandability) often result in a language user going through several (abortive) attempts at conveying their ideas. For instance, they may first have in mind one strategy for explaining what they want to say, and then switch to a different strategy because they think it will be more effective rhetorically.[75] I do not consider this sort of adjustment to be part of the process of expressing a single thought: rather, each change in rhetorical strategy involves a reformulation of the thought; and so each strategy requires a new attempt to put a different thought into words. Different strategies can overlap, however, so each new attempt does not necessarily start from scratch. Also, in general, when speaking before the entire expression is decided, the expression of the remaining (i.e. not yet expressed) part of the thought is built around what has already been said, unless it is 'cancelled' and re-done from the beginning (see Section 5.3 for more discussion). This can apply to changes in pragmatic strategy as well. So even if the thought to be expressed changes, (some) already-expressed linguistic material can be preserved.

As a side note, that I universally use the term *expression*, and not (in its place) *sentence* or *utterance*, is a reflection of some of these concerns: the expression does not need to be a sentence, and it does not need to be uttered.

With these points recognized, we can now discuss how the schema of the process of expressing thought can be described.

We know where our problem starts: it begins with a thought. Whatever is turned into an expression must first exist as a thought. Since thought is not conducted in linguistic expressions, it needs to be 'translated' into language. The thought will have a particular conceptual structure.

We know where it ends: with a linguistic expression. The linguistic expression will depict a particular conceptual structure. Linguistic expressions are arrangements of linguistic elements; the meanings of each component element and the overall meaning of the expression are peculiarities of the language. So the expression will be an approximation of the thought.

Our focus is on what goes on between these representations.

[75] See Gumperz (1982), Levelt (1989), and Carston (2002) for detailed discussions of the role of pragmatics.

Some models divide the process into semantic encoding and syntactic encoding (e.g. Levelt 1989; Levelt et al. 1999; Pickering and Garrod 2013[76]) (c.f. the discussion on architecture above). My architecture has syntax and semantics in a single level, so there are no separate steps regarding this. Manipulating syntactic and lexical structures is very much an integrated endeavor. As Bock and Ferreira (2014) discuss, building expression around a word or collection of words requires manipulating syntax, and building an expression around syntactic structures requires manipulating words.

In order to construct an arrangement of elements to express a thought, a language-user needs to select the linguistic elements that are to comprise the arrangement and put them together (these are facets of the same process, not different processes). In order to do this, the information to be expressed must be organized into available linguistic elements. This is discussed in Chapter 4.

Since the conceptual structure depicted by the expression is non-identical to that of the original thought, the process of expressing thought in language involves (re-) conceptualization. The idea that expressing thought in language is related to conceptualization has been frequently noted by other researchers (e.g. Levelt 1989; Levelt et al. 1999; Croft 2007; Langacker 2008). For example, the same physical situation can be described linguistically as *I see a chair with a ball to the left of it* or *I see a chair with a ball to the right of it* (Levelt et al. 1999: 8-9); expressing the idea in either of these forms is resultant from conceptualizing the scenario in a particular

[76] Pickering and Garrod state that they are assuming that semantics is constructed before syntax (without discussing a justification for that assumption), but allow that syntax can be decided first, as a sort of "prediction". They explain: "For example, a speaker might decide to describe a transitive event. At this point, she constructs a forward model of syntax, say $[NP\ [V\ NP]_{VP}]_S$, where *NP* refers to a noun phrase, *V* a verb, *VP* a verb phrase, and *S* a sentence. This forward model appears appropriate if the speaker knows that transitive events are usually described by transitive constructions, a piece of information assumed in construction grammar (Goldberg 1995), which associates constructions with "general" meanings. The speaker can therefore make this prediction before having decided on other aspects of the semantics of the utterance, thus allowing the syntactic prediction to be ready before the implemented semantics." (p. 339). This characterization is in line with the theory I advance here, except I see no reason to discriminate between syntax and semantics in the first place.

manner. This is true of every aspect of creating a linguistic expression, since linguistic meaning *is* conceptualization—conceptualization is what dictates use of a particular word or construction, and conceptualization is necessary for creating the mental representation of the expression.

A point of caution should be noted, however. In cognitive linguistics, conceptualization or construal is given great emphasis (and properly so). With regard to the topic of the expression of thought in language (which is not a major one in cognitive linguistics; it is generally mentioned only in passing, and rarely at that), the explanation given is entirely in terms of construal. The following two statements are representative:

> "[C]onstrual is a means by which various steps in the verbalization process are taken." (Croft 2007: 355)

> "At the conceptual level, we are presumably able to evoke…content in a fairly neutral manner. But as soon as we encode it linguistically, we necessarily impose a certain construal." (Langacker 2008: 43)

Now, while I do not dispute this understanding—rather, I endorse it—, the sort of one-sided characterization typical in cognitive linguistics of construal as a property of the conceptual structure of linguistic expressions (to the neglect of that in language-independent thought) can easily lapse into the contention that there is no conceptual structure until thought is put into language, or that construal exhaustively accounts for

the process of expressing thought in language[77]. The moral is: do not conflate conceptualization as part of the process of expression of thought with the 'natural' conceptualization that is part of the formation of the thought.

Both the organization of conceptual information and building the conceptual structure that is the meaning of the expression are related to meaning construction. The principles of meaning construction are what make an arrangement mean what it does. This is discussed in Chapter 3.

So, our task in building a theory of the expression of thought in language is to describe the process of taking a language-independent representation (the original thought) and deriving an approximately-equal language-dependent representation (the expression) from it (Figure 6).

[77] A moment's reflection will show that (re-)construal alone cannot effect the expression of thought: causality is reversed. Since linguistic expressions approximate the conceptualization in the thought, the construal in the original thought is a major driving force that shapes the expression. For example, one says that a cup is half-full or half-empty because first one thinks that it is half-full or half-empty. But the conceptual structure actually depicted by an expression depends on the particularities of the language, and this cannot be dictated absolutely by the original construal. For instance, the closest Japanese equivalent to *half-full* is *hanbun haitteiru* (lit. 'has half gone in'); the conceptual structures of the Japanese and English expressions are not the same (it is true in general that conceptual structures in different expressions (or languages) are different, though the differences are not so striking in some cases). So expressing an idea in a particular way because that is how it is construed is only half the story; the other half is that it needs to be construed in a particular way in order to express it that particular way (c.f. Slobin 1996).

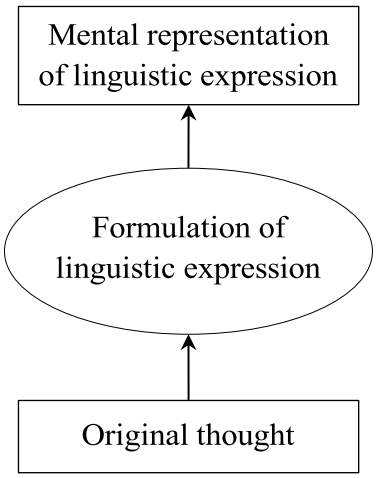

Figure 6: The problem in basic outline

We can schematize this process as

$$\Lambda_\alpha \Rightarrow \Lambda_\beta;\ \Lambda_\alpha \approx \Lambda_\beta,$$

where Λ_α is the original language-independent thought (the starting representation) and Λ_β is the linguistic expression (the product representation derived from it). The expression is comprised of an arrangement of linguistic elements; if we denote the arrangement of elements as $\rho_1 + \ldots + \rho_n$, we can write this as

$$\Lambda_\beta = \rho_1 + \ldots + \rho_n.$$

Finally, the arrangement of elements that is the expression yields an approximation of Λ_α according to the mechanisms of meaning construction. That is:

$$\rho_1 + \ldots + \rho_n \rightarrow \Lambda_\alpha.$$

This schematization will be relevant in Chapter 5.

So this is our characterization of the problem. The central claim of this dissertation is that the cognitive process responsible for creating a suitable expression is a kind of abduction driven by background knowledge of meaning construction (Chapter 5). The following two chapters set the discussion of abduction up by examining meaning construction (Chapter 3) and constructing of arrangements, which consists of organizing the conceptual information to be expressed and selecting the linguistic elements of the expression (Chapter 4).

Chapter 3

Meaning Construction

Using an arrangement of linguistic elements to express a thought relies very fundamentally on that arrangement having a particular meaning. The means by which that comes to be so is called *meaning construction*; this is the subject of this chapter. A fact crucial for our explanation of the cognition of expressing thought in language is that a particular arrangement of linguistic elements can be relied on to yield a particular meaning. That is: the meaning of an expression is necessarily predictable. My main goal in this chapter is to characterize meaning construction as a dynamic phenomenon—where the yielded meaning is not simply a structured list of component meanings—that is based on an understanding of how meanings function and interact, which serves as a background theory on the basis of which the constructed meanings are predictable.

The predictability of the effective meaning of linguistic expressions is known as *compositionality*. In Section 3.1, I clarify what is meant by my assertion that the meaning of an expression is necessarily predictable. There are certain ways in which compositionality is (or, rather, can be seen as being) 'messy'; I discuss how these are not problematic for that assertion. In Section 3.2, I describe how meaning construction is dynamic. I also note that meaning construction applies both to combinations of meanings and to individual meanings in isolation. In Section 3.3, I turn to the issue of the cognitive basis of meaning construction. I propose that the conceptual understanding of how meanings function and interact serves as a background theory that allows language users to predictively 'infer' what meaning will be yielded by a particular arrangement. I discuss how the cognition involved in this understanding is related to mental model cognition and simulation cognition. I summarize the chapter in Section 3.4.

3.1 ‘Degeneracy’ in compositionality

It is a necessary fact that language is compositional; otherwise it could not be used. The fact that people can understand and produce expressions without having individually memorized them in all of their infinity depends on compositionality[1]. However, despite this necessity, there are some apparent imperfections in the compositionality of language.

Some linguistic structures appear to be composed of parts and yet not have their meaning be determined by the meanings of those parts. Examples of this include idioms like *kick the bucket* (where kicking and buckets have nothing to do with the overall meaning), phrasal verbs like *look up* (where ‘look’ and ‘up’ do not determine the overall meaning), and historical compounds like *understand* (where the (current) meaning is not built from ‘under’ and ‘stand’; Langacker 2008: 170). In these cases, it appears that the meaning of the whole is not determined from the parts.

The issue here is that not every bit of linguistic structure that is conceivably a component actually functions compositionally. This is one of the points I brought up previously in Section 2.2.1.2 (What counts as an element?). The same issue is seen in *squealer* vs. *drawer* (where *squealer* is indeed built from <squeal> + <-er>, whereas *drawer* is not built from <draw> + <-er>; Langacker 2017: 43 etc.), *replicate* (which is generally not analyzable into <re> + <plicare>; Levelt et al. 1999: 12), and *metaphor* (which may or may not be identified as being formed from two morphemes (<meta> and <phor>), whose meanings (‘across’ and ‘carry’, respectively) may or may not be

[1] Some articulations of the notion of compositionality (e.g. “The meaning of a whole is a function of the meanings of the parts and of the way they are syntactically combined.” (Partee 2004: 146)) presume a distinction between the lexicon and syntax. Others do not; e.g. Langacker (2008) states: “An expression is said to be compositional to the extent that its composite structure derives in a regular, predictable way from its component structures. Compositionality is an essential feature of language, enabling us to create and understand an endless supply of new expressions.” (pp. 167-168). I, like Langacker, include syntactic combinations as part of the component parts of expressions whose meanings determine the overall meaning, as discussed in Chapter 2. There are other variations in the details of different conceptions of compositionality; see Pelletier (2017) for discussion. I am using the term to refer specifically to the systematic predictability of the meanings of arrangements of linguistic elements.

known, and which may or may not be understood as being formed from those meanings). As noted in that discussion, not every bit of linguistic structure that might conceivably be a linguistic element necessarily counts as one. Linguistic elements do not necessarily correspond in a one-to-one fashion with every part of linguistic structure that might be identified as a 'unit'.

In this sense, compositionality does not inhere at the level of the bits of linguistic structure that might be identified as units, but rather at the level of actual linguistic elements. The pieces of linguistic structure that do not count as linguistic elements do not individually participate in meaning construction.[2]

It should be stressed (again) that there is a gradience in mismatch between the 'formal' structure from which one might demand compositionality and find it lacking and the elements that actually participate in meaning construction and which do exhibit compositionality. This depends also on the analysis of the individual language user; c.f. also footnote 72 in Section 2.3 (The architecture of language). One clear example is the idiom used in *Suzie sang her head off* ('Suzie sang a lot/intensely'; Jackendoff 2017: 194). In this case, Suzie's head does not literally come off, of course. There are two basic analyses possible here. One is that the expression <VERB one's head off>[3] has an entrenched figurative meaning that does not involve any imagery of her head popping

[2] Cognitively, the import of this fact is not restricted to compositionality itself. In Section 2.2.1.2, I gave the example of *in order to* as potentially (though not necessarily) being a linguistic element (i.e. directly associated with a particular meaning and accessible as a whole without having to be accessed through its component parts); *in order to* has no 'problems' with compositionality, in the sense that the compositional meaning derived from the combination of the component meanings is entirely compatible with the overall meaning. But if *in order to* as a whole is a linguistic element (for a particular language user), then its meaning is not actually composed from the component meanings. So in this case, there would be no actual compositionality (as a cognitive matter) at the level of the component bits of *in order to*, despite the fact that they can be identified as units and, from an 'objective' theoretical standpoint, they seem to generate the correct overall meaning.

[3] As mentioned in Section 2.2.1.3 (Granularity of linguistic meaning), the generality of constructions is often not fixed; it is possible, for instance, to take this construction to not require specifically the head (e.g. it is possible to consider *The pianist played his fingers off* to employ the same construction).

off as a result of the intensity of her singing. In this case, the overall meaning is not compositional from the components that do mean that her head came off. The other is that it *does* involve such imagery, and arrives at the non-literal meaning through hyperbole. In this case, it is compositional from the components that mean her head came off. In the first analysis, the 'formal' structure does not have compositionality, and in the second analysis, it does.

So these cases do not violate compositionality; there is no difficulty in understanding what they mean or producing them to express a particular desired meaning. What is instead involved is variation in what counts as an element.

Another apparent sort of degeneracy is seen in cases where the 'compositional' meaning is underdetermining of the presumable 'actual' meaning.

One example is the three expressions *The child is safe*, *The beach is safe*, and *The shovel is safe*: these all can, for instance, be said in the context of a child playing with a shovel on a beach; and in this context, they can all refer to the safety of the child (Fauconnier and Turner 2002: 25). The specific differences between them—the details of how the child, the beach, and the shovel are safe, and for whom—are not reflected in the form of the expression.

Similarly, noun compounds have a variety of relationships between the nouns, which are not made explicit in the form of the expression. Jackendoff describes the full interpretation of expressions like these as involving "extra semantic material" not present in the expressions themselves; in the examples shown below provided by him (2017: 204), this material is indicated with underlining.

beef stew = ‘stew made of beef’
attack helicopter = ‘helicopter that attacks’[4]
mosquito netting = ‘netting to protect from mosquitoes’
dog house = ‘house for a dog to live in’
bike helmet = ‘helmet to wear while riding a bike’
paper clip = ‘clip to hold papers together’
oil stove = ‘stove that runs on oil’

Langacker (2008: 168) gives the examples of *toothbrush*, *alarm clock*, *pear tree*, *peanut butter*, *tablespoon* to the same effect: the construction does not specify exactly how they are related, beyond that a *toothbrush* is a particular kind of brush that has to do with teeth, etc. Langacker also notes (ibid.: 169) that *jar lid*, while it generally means ‘the lid of a jar’, can have other meanings, like ‘a lid decorated with a picture of a jar’.

Regarding phenomena like these, there are two possibilities concerning the relevant mental representation. Which of them applies again depends on the individual language user’s analysis, of course.

i) The actual meaning is schematic; while a more specific interpretation may be true of the content in question (e.g. it is true that a dog house is a house for a dog to live in), that specific interpretation is not actually part of the meaning (e.g. the actual, schematic meaning can be interpreted as either ‘the lid of a jar’ or ‘a lid

[4] “Properly” speaking, *attack* here is a technical term referring to a particular combat role (aircraft engaging individual targets on the ground); in this sense, *fighter aircraft* (which specifically engage enemy aircraft, and are not just ‘aircraft that fight’) are distinguished from *attack aircraft*. Ultimately, however, what is relevant for us is what is actually represented in the mind of the language user; this is a matter of individual analysis. If Jackendoff or any other language user actually understands the expression as meaning ‘helicopter that attacks’, that is reality as far as we are concerned. Jackendoff’s “mistaken” gloss is actually an excellent demonstration of this point.

decorated with a picture of a jar' etc., but it itself does not specify from among them) (this is Langacker's (2008: 167-170) explanation[5]).

ii) Each specific 'interpretation' (to the extent that it is cognitively real) corresponds to a distinct linguistic structure; so the forms corresponding to those structures are just ambiguous.

This is the same problem discussed in Section 2.2.1.3 (Granularity of linguistic meaning): the absence of a fixed delineation between schematicity and polysemy. It is important to note (again) that if the language user can intend a specific meaning in their original thought, but the specificity in the thought is not necessarily reflected in the (mental representation of) the linguistic expression, so i) can hold even if originally the intended meaning was more specific.

At any rate, in either of these cases, compositionality still holds. In the first case, the overall meaning is fully determined by the component meanings; it is just a schematic meaning. In the second case, the overall meaning is likewise fully determined by the component meanings; there is just ambiguity from several meanings being associated with the same form.

While the issues discussed in this section—that what actually function as linguistic elements do not necessarily match up with what might be discerned from the formal linguistic structure and apparent underdetermination of meaning in expressions—do make theoretical description of language less straightforward, they do

[5] Langacker actually says with regard to these multiple possibilities for *jar lid* that "[*j*]*ar lid* is arguably not quite fully compositional because the constructional schema does not guarantee that the association between jar and lid will be the obvious one of the lid serving as cover for the jar" (2008: 169 note 8). He then gives examples like *understand* (mentioned above), adds that general knowledge, understanding of context, and devices of conceptualization like metaphor and fictivity (discussed in Sections 2.1.4.1 and 2.2.1.4) also contribute to establishing the meaning of expressions, and concludes that "most expressions are only partially compositional" (pp. 169-170). This point relates to the topics of Sections 3.2 and Section 3.3 (essentially, it falls into what I am referring to as dynamicity, and the conceptualization involved in meaning construction is part of what I refer to as mental model cognition), but I would reject (as I explain in the paragraphs immediately below) the notion that this makes language less than wholly compositional in the sense that I mean.

not present any obstacle to the predictability of meaning. It is still absolutely true that a particular arrangement of linguistic elements can be counted on to yield a particular meaning.

3.2 The dynamicity of meaning construction

Now, that this is true does not mean that the meaning of an expression is just a structured list[6] of the component meanings. The point I wish to make in this section is that meaning construction is produced by the particularities of the meanings of elements and how they interact with each other. The sort of synthesis of meaning that this involves is what I refer to as 'dynamicity'. In this section, I describe some of the variety in paths to meaning by which different meanings can be made. My 'argument' in this section, if there is one, is that entirely different paths approximating the same meaning is indicative of the sort of dynamicity that I am referring to. The meanings yielded by arrangements of elements are governed by dynamic principles.

---

[6] Structure of course needs to be taken into account. *Unstructured* lists are even more obviously not applicable to the meaning of expressions. Langacker (2008) explains:
"The simplest hypothesis would merely identify an expression's composite meaning with the set of its component meanings. Composition would then be just a matter of viewing the component meanings collectively. On this account, the composite meaning of *jar lid factory* would be the unordered set {[FACTORY], [LID], [JAR]}. It is readily seen, however, that there is more to composition than mere summation. Otherwise, distinct expressions with the same components would always be semantically equivalent. But they are not. We cannot, for instance, ignore the semantic differences of *jar lid factory* (factory for making jar lids), *lid factory jar* (jar used in a lid factory), and *jar factory lid* (cover for a roofless jar factory). An expression's composite meaning is not just a pile of component meanings, but an integrated structure where elements relate to one another in very specific ways." (p.168)
My point here is that even *structured* lists cannot fully capture the meaning of expressions.

Of course, in some cases, the meaning of an expression is actually a list (e.g. *apples, mangos, and pasta* or *I need a pen, a desk, and a chair* or *He is undisciplined and undriven*). But these are special cases where the construction involved specifically indicates that meaning (that there is a list); the 'listness' does not simply emerge naturally from an agglomeration of multiple elements.

A clear demonstration of this fact is the great variety in 'paths' taken by linguistic expressions to express meaning. Similar effective meanings can be gotten to from wildly different 'starting points', and a particular bit of meaning can contribute to wildly different effective meanings. The mechanisms of meaning construction—the 'dynamic principles' referred to above—are what makes a path lead to its destination. Entirely different arrangements of elements can express approximately the same meaning; some examples are shown below.

(1) a. behold (English)
b. hineh (Hebrew) (lit. 'here') 'behold'
c. thoroughly hold (etymological meaning of *behold*)

(2) a. (and) also (English)
b. necnon (Latin) (lit. 'and not not') 'and also'

(3) a. necessary (English)
b. hitsuyou (Japanese) (lit. 'absolute importance')
c. necessus (Latin) (lit. 'not yielded')

In (1), we see three different ways of getting to a meaning like 'behold': one is the English (1a), which gets there directly; another is the Hebrew *hineh*, which literally means 'here', but through the sort of dynamic meaning construction I am talking about can effectively mean something like 'behold'; and the third is 'thoroughly hold', which is the etymological meaning of *behold*. These are different paths that get to the "same" meaning. In (2), we see that 'and also' (2a) can also be gotten at by 'and not not' in Latin (2b); this is another path to the same effective meaning. Likewise, in (3), we see different ways of getting to 'necessary': in addition to the English (3a), it can be obtained by something like 'absolute importance' in Japanese (3b) or 'not yielded' in Latin (3c). Again, there are several ways to get to the same effective meaning.

The same phenomena underlie semantic extensions that happen as part of historical linguistic change. In Yoruba (a West African language), the verb *fún* has a dative function (shown in example (4) below), which developed from the meaning 'give' (Croft 2003: 34).

(4) wọ́n á sọnwó fún mi
they will pay.money DAT me
'They will pay me.' (Croft 2003: 34, citing Bamgbose 1966: 77)

A similar extension is found in Chinese *gěi*. As in (5a), it can express giving directly. It can also have a dative-like role where it does not specifically mean 'give'. In (5b), *gěi*'s function is quite analogous to that of *fún* in (4) above, where it indicates the recipient of the paying. It should be noted that this meaning of *gěi* is quite abstracted away from 'give', to the point where it can be used in situations that do not involve transfer to the recipient at all; this is well-illustrated in (5c), where the object in question is being *taken away from* the "recipient" marked by *gěi*.

(5) a. tāmen (yào) gěi wǒ yí-fèn lǐwù
they (will) *gěi* 1 one-CLF present
'They will give me a present'
b. tāmen (yào) gěi wǒ fù-qián
they (will) *gěi* 1 pay-money
'They will pay me'
c. kěyǐ gěi wǒ ná zhège ma?
can *gěi* 1 take this Q
'Can you take this for me?'

Here, the path to meaning goes from something like 'give' to something like 'to' or 'for'. What makes this possible is that 'to/for X' can be considered to be an abstract sense of 'give X'. This abstract sense of 'giving' applies even when there is no physical giving, as in (5c), for example; the link there is that doing something as a favor for me constitutes giving me a thing figuratively. Understanding how the meaning of something like 'give X' functions is crucial for establishing this abstract sense. It should be noted as well that English, for instance, takes a different path to the equivalent meanings: it uses *to* or *for* and their associated constructions or the <VERB X NP> construction (which all have their own nuances), none of which are related to 'give'.

Similarly, in Tzutujil (a Mayan language), the meaning of *majik* (‘because of/on account of’) in example (6) developed from the meaning ‘sin’ (Croft 2003: 34, citing Dayley 1985: 153).

(6) Xch’ejyi jar iixoq ruu- majik jar aachi
was.hit the woman 3SG.POSS- because.of the man
‘The woman was hit because of the man.’

It is a general truth that there are many ways of expressing a particular sort of meaning. I will discuss a few examples below, but it should be kept in mind that the examples brought up here represent only a tiny part of the overall diversity seen in language.

In the English *He ran fast*, the method of attaching ‘fast’ to ‘ran’ may seem a matter of course. But this is not the only way of accomplishing it. In the following examples, a similar construction for attaching the description of manner (‘fast’) to the verb (‘run’/‘go’) is used in Hausa (a West African language) (7a) and Classical Mongolian (7b) (Croft 2003: 11-12). They both achieve this by expressing it as ‘run with speed’. So one pattern here constructs the meaning as ‘with speed’, and the other does so as ‘speedily’. Hebrew expresses this idea using a construction quite similar to that used in English (‘run fast’) (7c), but also has a construction analogous to those used in the Hausa and Classical Mongolian in (7a) and (7b); however, this yields a somewhat different meaning (7d). It is not the overall form of the expression that determines the sort of meaning that is constructed, but rather the particularities of the meanings of the individual elements involved (which are slightly different) and how they interact that do.

(7) a. yā gudù dà saurī
3SG.COMP run with speed
‘He ran fast (“with speed”).’
(Croft 2003: 11, citing Kraft and Kirk-Greene 1973: 85)

b. türgen -iyer yabumui
speed -with goes
‘He goes fast.’ (Croft 2003: 12, citing Poppe 1974: 153-154)

c. hu ratz maher
   he ran fast
   'He ran fast.'
d. hu ratz be-mehirut
   he ran with-speed
   'He ran hurriedly'

One way of getting at the meaning of 'sometimes A; sometimes B' is to alternately focus on the different states. Zaza (an Iranian language spoken in Turkey) (8a) and Russian (8b) use similar constructions to this effect (Haspelmath 2007: 27). Whereas the *sometimes* used in the English construction is indefinite, Zaza and Russian use 'now' (which is not indefinite—it is specific, although freely-set) in the following examples.

(8) a. Na zu __ hewro, __ pakao
   these days now cloudy now clear
   'These days, sometimes it is cloudy; sometimes it is clear'
   (Haspelmath 2007: 27, citing Selcan 1998: 667)

b. Xolodnyj do to usilivalsja, to oslabeval
   cold rain now strengthened now weakened
   'The cold rain sometimes became strengthened, and sometimes weakened' (Haspelmath 2007: 27)

It should also be noted that this sort of expression involves a shifting perspective (c.f. Section 2.2.1.4.2): the perspective moves from points where it is cloudy to those where it is clear in (8a), and from points where the rain is stronger to those where it is weaker in (8b). By contrast, the English sometimes it is cloudy; sometimes it is clear and the cold rain sometimes became strengthened, and sometimes weakened do not have a shifting perspective.

A structurally similar construction, with a different meaning, exists in Chinese. The <yībiān A yībiān B> construction expresses that A and B are done simultaneously, in parallel. Two examples of it are shown in (9) below.

(9) a. wǒ xǐhuan yībiān hē chéng zhī yībiān dú bàozhǐ
I like *yībiān* drink orange juice *yībiān* read newspaper
'I like to drink orange juice while I read the newspaper.'

b. tā bù něng yībiān xiǎng shìqing yībiān hūxī
he not can *yībiān* think thing *yībiān* breathe
'He can't think about something and breathe at the same time.'

A comparison of the Chinese expressions and their English translations shows different ways of effecting similar meanings: the relationship of parallel simultaneity is expressed through different paths.

On the same general topic of the relationship between different events, some interesting manners of constructing meaning can be found in Tübatulabal (a Californian Uto-Aztecan language). In Tübatulabal, morphological marking on verbs in a subordinate clause[7] distinguishes between anterior action and two kinds of simultaneous action: in the case of anterior action, one event is completed when the second occurs (10a); in the case of one kind of simultaneous action, the first event continues through the occurrence of the second (10b); and in the case of the other kind of simultaneous action, the first event is interrupted by the occurrence of the second (Corbett 2007: 309, citing Voegelin 1935: 126-127).

(10) a. kó:imí ánaŋ-í:yá'awáŋ iŋgím tá:twál
woman cry-ANT came man
'When the woman had stopped crying^{ANT}, the man came.' (Corbett 2007: 309)

b. kó:imí ánaŋ-áŋ iŋgím tá:twál
woman cry-SIM1 came man
'While the woman was crying^{SIM1}, the man came.' (ibid.)

[7] Note also that neither event is structurally subordinate to the other in the Zaza, Russian, and Chinese constructions discussed above in examples (8) and (9).

c. kó:imí tïka-káŋ apá'agín tá:twál
woman eat-SIM2 hit man
'The man hit the woman when the woman was eating^{SIM2} [and as a result her eating was interrupted].'[8] (ibid.)

These are tools available to express meaning in Tübatulabal, and they enable rather different paths to meaning from those possible in English and many other languages, which do not have any similar elements[9].

Similarly, there are different ways of getting at the sort of meaning expressed by *The king eats whatever he wants* (11a). The Chinese in (11b) 'zooms in' to individual cases of wanting: If the king wants to eat some thing, he eats that thing.[10] The overall structure is rather different. The Chinese employs a conditional, whereas there is no conditional in the English. The structural relationships in the conception depicted (c.f. Section 2.2.1.4.4) are different as well: the English depicts only one object, which is both wanted and eaten (*whatever*), but the Chinese depicts two objects (*shénme… shénme*), one of which is wanted and the other of which is eaten, that are equated as the same thing. There is no particular element in (11b) that corresponds to *whatever* in (11a); *shénme* by itself does not have approximately the same meaning as *whatever*, and there are two of them. Rather, the equivalences emerge at a holistic level of meaning construction. Despite all of these differences, these arrangements yield effectively approximately-equal meanings.

(11) a. The king eats whatever he wants.
b. guówáng xiǎng chī shénme, jiù chī shénme.
king want eat what then eat what
(lit. '(the) king wants to eat something, then (he) eats something')

[8] This seems to be a popular theme.

[9] Issues relating to differences in meaning between these Tübatulabal expressions and English near-equivalents using *while* and *until* will be discussed in Section 4.1.2.2 (Interdependencies).

[10] It should be noted that it is possible to analyze *whatever he wants* as being directly generic or as applying firstly to some individual instance of wanting (i.e. 'whatever he wants in that specific instance') and subsequently being generalized to all such instances. The second analysis therein has similarity to the Chinese expression.

One reason for the variety between languages in paths to meaning is that the elements that each language has to work with are different, even when they appear similar. As a result of this, often, elements in different languages will seem to correspond to each other nicely in some linguistic contexts, but in different contexts, a separation is seen. For instance, Croft (2003: 6-7) gives the following examples of cases of alignment and misalignment of article use in English and French. In (12), they align: where one language uses a definite article ('the') or an indefinite article ('a') or no article at all, so does the other. Just looking at these cases, one might conclude that the definiteness systems in French and English are essentially the same.

(12) a. He broke a/the vase.
Il a cassé un/le vase.
he has broken a/the vase

b. The concert will be on Saturday.
Le concert será Samedi.
the concert will.be Saturday

c. He went to the bank.
Il est allé à la banque.
he is gone to the bank

However, in other cases, such as those in (13), there is no such neat correspondence. Here, where one language uses no article, the other does use one. Additionally, here we see that articles in French contain information about number and gender, which is absent in the English article system (in the French article system, there is no plural indefinite article—although, as shown below in (13e) and (13g), *des* can have that sort of function—(so distinctions in plurality only exist for the definite articles), and distinctions in gender exist only for the singular form (the feminine and masculine definite articles—*la* and *le*, respectively, collapse to the same plural form, *les*)[11]. Also, French uses expressions like 'of the'[12] in ways that are rather different from

[11] In French, adjectives and nouns also have information about number and gender (see Section 4.1 for more discussion), but in the gloss below, I have only indicated that information for the articles).

the usage in English; that this is possible is due to the fact that 'the' (*le*/*la*) in French can have a different meaning from *the* in English.

(13) a. I drank wine.
J'ai bu du vin.
I have drunk of.the.M wine

b. The French love glory.
Les Français aiment la gloire.
the French love the.F.SG glory

c. He showed extreme care.
Il montra un soin extrême.
he showed a.M care extreme

d. I love artichokes and asparagus.
J'aime les artichauts et les asperges.
I-love the.PL artichokes and the.PL asparaguses

e. Birds have wings.
Les oiseaux ont des ailes.
the.pl birds have of.the.PL wings

f. His brother became a soldier.
Son frère est devenu soldat.
his brother is become soldier

g. Dogs were playing in the yard.
Des chiens jouaient dans le jardin.
of.the.PL dogs were.playing in the yard

A similar sort of partial alignment can be seen in the definite article distributions of English and Hebrew. In some cases, they correspond rather well (14).

[12] In the examples below, only *du*, contracted from *de* ('of') + *le*, and *des*, contracted from *de* + *les*, are included; *de la* (which does not get contracted) is missing.

(14) a. ha-ra'ayon she-hitzati
the-idea that-I.suggested
'the idea that I suggested' (adapted from Glinert 1989: 13)

b. ha-mechonit be-tikun
the-car in-repair
'The car's being fixed' (Glinert 1989: 14)

c. otam ha-shemot
same the-names
'the same names' (ibid.: 97)

d. aleksander ha-gadol
Alexander the-great
'Alexander the Great'

But in other cases, like those in (15), they do not. There is no 'the all' (15a) or 'the these' (15b) in English. Again, that the usages in the two languages align only partially is because the meanings of the corresponding elements are different.

(15) a. ani ochel ha-kol
I eat the-all
'I eat anything' (Glinert 1989: 13)

b. ele ha-dapim
these the-pages
'these pages' (ibid.: 97)

There is no segregation of systems of meaning construction according to the size or level of the relevant bit of the expression. Meaning construction involving sub-lexical morphology works in the same way as meaning construction at the phrase level. Bruening (2018) gives numerous examples of words formed with phrasal syntax, like *She had that I'm-so-proud-of-myself look*. (p. 3), which shows that exactly the same sort of meaning construction can be involved in linguistic structures larger than words and in those smaller than words. It should be noted that this unity of mechanism applies to 'normal' polymorphemic assembly of words as well. The basic mechanisms of meaning

construction that make *unthinkable* in English or *doutou* (lit. 'same-equal'; 'equivalent') in Japanese mean what they do are not separate from those involved in the construction of meaning at the phrasal level. Any differences between sub-lexical-level meaning construction and supra-lexical meaning construction lie only in what the particular linguistic elements involved are, and not in any dissimilarity in *how* those elements are used to build meaning.

Another point concerning the sort of dynamism in meaning construction that I am talking about that should be noted is that particular concepts can be gotten at with composite linguistic representations—ones comprised of multiple component bits of meaning—or with indivisibly 'atomic' linguistic representations[13]. Some examples are shown below in (16)-(22); of each pair in these examples, one is generally analyzable only as an atomic object, and the other is generally analyzable as a composite object. Of course, whether a particular linguistic structure is actually (in terms of cognitive reality) composite or atomic can depend on the analysis of the individual language user (c.f. Section 2.2.1.2 (What counts as an element?)).

(16) a. some (English)
b. nonnulli (Latin) (lit. 'not none')

(17) a. <u>likely</u> outcomes (English)
b. <u>kanousei ga takai</u> kekka (Japanese) (lit. 'likelihood-is-high outcome(s)')

(18) a. extreme (English)
b. hijou (Japanese) (lit. 'not normal')

(19) a. kuzureru (Japanese)
b. fall apart (English)

[13] To clarify my wording in this sentence: I am using the word *concept* loosely here. Strictly speaking, it should be recognized that two different linguistic representations of the "same" concept are in fact different concepts. Also, I do not want to give the impression that any bunch of material corresponding to what, in the linguistic expression, can be identified as a "concept" necessarily exists as a unit in the original thought. If one looks at a linguistic expression, it may seem that certain chunks of it form 'natural concepts', but I do not intend to assert or imply that those 'natural concepts' must correspond in a one-to-one manner to fundamental units in the original language-independent thought (c.f. Chapter 4).

(20) a. umbrella (English) (itself from ‘little shadow’, but generally not analyzable thusly)
b. paraguas (lit. ‘stop waters’) (Spanish)

(21) a. length (English)
b. cháng-duān (Chinese) (lit. ‘long-short’) (Aikhenvald 2007: 25)

(22) a. size (English)
b. dà-xiǎo (Chinese) (lit. ‘big-small’) (ibid.)

Meaning construction is quite fluid between atomic and composite representations. Approximately-equal meanings in expressions can be yielded by atoms and other atoms (e.g. *rage* vs. *fury*), atoms and composites (as in (16)-(22) above), and composites and other composites (e.g. *hitsuyou* (‘absolute importance’) vs. *necessus* (‘not yielded); ((3) above).

Another point that should be recognized is that meaning construction is not limited to combinatoriality. Much of the importance of meaning construction is in combining multiple linguistic elements to yield some effective meaning. But meaning construction is also important in arrangements or parts of arrangements that consist of single elements. For example, in the expressions shown below, there is noticeable meaning construction involved in deriving effective meanings from single elements. Again, this linguistic process is responsible for diachronic semantic extension of or changes in lexical meaning. In (23a), a meaning like ‘even’ is constructed from the single-element meaning ‘arrived’; from the meaning ‘here’, effective meanings like ‘behold’ (23b) or ‘I’m giving you this’ (24c) can be constructed; similar phenomena are seen in other examples that were discussed above (23d-f).

(23) a. La semana pasada, mi madre fue muy mandona. Llegó a decirme que limpie mi cuarto.
(lit. ‘Last week, my mother was very bossy. She arrived at telling me to clean my room.’ = ‘…She even told me to clean my room.’)
b. hineh (‘here’ → ‘behold’) (Hebrew); c.f. (1b)
c. here → ‘I’m giving you this’ (English)
d. fún (‘give’ → ‘to’) (Yoruba); c.f. (4)

e. gěi (‘give’ → ‘to’) (Chinese); c.f. (5)

f. majik (‘sin’ → ‘because of’) (Tzutujil); c.f. (6)

Now, as I remarked in footnote 74 in Chapter 2, it is possible for single-word expressions (like *Food!*) to involve constructions as well. This does not take away from the point that in cases like those in (23), the action of meaning construction is not combinatorial—even if there is a construction that the word is participating in, the construction does not contribute to these aspects of the effective meaning, so the word is not being combined with the construction to build these aspects of meaning[14].

3.3 The cognitive basis of meaning construction

The basic point made in the previous section is that meaning construction is grounded in dynamic principles that govern what meanings are yielded by particular arrangements of elements; language users must know these principles in order to use language. In this section, we address the question of the sort of cognition that underlies meaning construction. I propose that understanding of how meanings function and interact works as a background theory on the basis of which the meanings of arrangements of elements are predictable. I have brought attention to the fact that constructed meaning is dynamic and not formed as a structured list of component meanings. Concerning this point, Langacker (2008) says “a complex expression’s meaning cannot be *computed* from lexical meanings and compositional patterns…but is more accurately seen as being *prompted* by them” (p. 245; emphasis original); I believe he is saying the same thing as I am when I say it is dynamic and not a structured list.

Other researchers have identified meaning construction as related to inference. I discuss this idea in Section 3.3.1. I adopt this view; however, ‘inference’ by itself does not amount to a detailed explanation. My main proposal is that meaning construction is

[14] Linguistic context can restrict what meanings are possible (so, from the standpoint of interpretation, it may seem that the construction etc. ‘makes’ a certain meaning be interpretable or not), but this is a different issue from the actual construction of meaning itself. See the discussion in Section 3.3.1 below for examples of this.

based on mental model cognition; this is what makes the meaning yielded by a particular arrangement be predictable to the language user. I discuss this in Section 3.3.2. A somewhat related proposal has been made by other researchers, who advance the idea that the cognition of meaning construction is simulation. The sort of mental model cognition I am talking about is related to simulation, though it is not identical with it. I discuss this in Section 3.3.3.

3.3.1 Inference

Many researchers have remarked that meaning construction is related to inference. For example, Jackendoff (1997) notes that inferring one statement from another depends on the constructed meanings of those statements. He points out that inferences between statements are based not on the form of the statements, but rather on their meaning. Certain statements either entail or do not entail others, regardless of any structural similarities or lack thereof. Thus, while the entailments in (24) are valid, those in (25)—despite using analogous structures—are not.

(24) a. Fritz is a cat. Therefore, Fritz is an animal. (Jackendoff 1997: 32)
b. Bill forced John to go. Therefore, John went. (ibid.)

(25) a. Fritz is a doll. #Therefore, Fritz is an animal. (ibid.)
b. Bill encouraged John to go. #Therefore, John went. (ibid.)

In order to infer one statement from another, the actual meaning of those statements must be understood; the processing of these kinds of inferences or entailments is based on an understanding of the meaning yielded by the arrangement of elements. It is not based on form itself; inferential relation of linguistically-expressed statements does not require shared linguistic material between the expressions. For example, (26b) can be inferred from (26a), but there is no linguistic material that is shared between the expressions that can constitute the basis for this inference.

(26) a. Bill killed Harry. (ibid.: 186)
b. Harry died. (ibid.)

Knowing that *Bill killed Harry* entails that Harry died depends, rather, on an understanding of the function and interaction of the meanings comprising the two statements; these are the 'dynamic principles' I have been referring to. So in this sense, how different linguistically-expressed statements are inferentially related is intimately connected to the principles of meaning construction that govern what effective meaning is yielded by an arrangement of linguistic elements.

Additionally, the interpretation of what meaning is being constructed in a particular arrangement can be described as a kind of inference. This is of central importance in theories of comprehension in communication, such as relevance theory (Sperber and Wilson 1995, etc.; Wilson and Sperber 2002, etc.), which describes interpretation of utterances as inference on the assumption that the information is relevant. As an example, Wilson and Sperber (2002: 262-267) describe the inferences involved in the interpretation of *He forgot to go to the bank*. The word *bank* might conceivably refer to a river bank or to a financial institution; resolving which meaning it has is of course necessary for interpreting the statement. They give the following discourse context:

Peter: Did John pay back the money he owed you?

Mary: No. He forgot to go to the bank.

Wilson and Sperber explain Peter's interpretation of Mary's reply as being effected by making inferences with the assumption that Mary's utterance will be optimally relevant to him; these inferences are based on encyclopedic knowledge. The expectation resultant from this assumption is that Mary's utterance will adequately explain why John has not repaid the money. On the basis of encyclopedic knowledge, Peter has access to the postulation that forgetting to go to the bank-as-financial-institution may make one unable to repay money that they owe. This allows him to infer that Mary's utterance means 'John forgot to go to the bank-as-financial-institution', since that would satisfy the assumption that her utterance is relevant. He can then make further inferences to enrich his interpretation: John was unable to repay Mary because he forgot to go to the bank; and John may repay her the next time he goes to the bank

(i.e. his failure to repay her is due only to the fact that he forgot to go to the bank—and not because he intends to abscond with the money or did not have money in his account, etc.—and if he had not forgotten, he would have repaid her). Wilson and Sperber also note that the specific interpretation of *bank* may be narrower or looser than 'financial institution'. It may be interpreted not generically, but in terms of the particular occasion at hand (i.e. narrower), as a particular bank or whatever bank Peter would go to in order to retrieve the money to pay Mary back; the World Bank, for instance, qualifies as a financial institution, but any encyclopedic knowledge that Peter has about the World Bank would not support his interpreting *bank* in a sense that includes it. At the same time, it may also be interpreted in a looser sense: for instance, an ATM in a supermarket is not a financial institution, but it is (again, on the basis of encyclopedic knowledge) an entity that would be suitable as a kind of *bank* in this context; so Peter may interpret *bank* as 'bank-or-cash-dispenser'; in fact, John may regularly get his money from an ATM not located in a bank, and in this case, if Peter interprets *bank* as excluding ATMs, Mary's utterance would not be an adequate explanation.

Another significant account of the use of inference in interpretation was advanced by Haviland and Clark (1974); their theory concerns the use of inference for contextually situating expressions so they can be interpreted. The idea is that interpretation of certain parts of expressions depends on antecedent information. Haviland and Clark reasoned that the interpretation would be easier and faster in cases where the part in question has a "direct antecedent", as in (27a) and (28a), compared to cases in which the antecedent is "indirect", as in (27b), (28b), and (29). In order to understand what *The beer* refers to in (27a), one must refer to the antecedent *some beer* in the previous sentence; here, *some beer* directly refers to the same thing as *the beer*, and so can be considered to be a direct antecedent. The same is true for *an alligator* and *the alligator* in (28b). By contrast, *the picnic supplies* does not mention the particular beer referred to by *The beer* in (27b), and *an alligator* does not refer to the particular alligator denoted by *the alligator* in (28b). Rather, to interpret the beer as being included among the picnic supplies and the alligator that was received as a present as

connected to the alligator that was wanted as a present[15] presumably requires inference not necessary for the interpretation of the direct-antecedent cases. The interpretation of cases like (29) presumably requires a similar sort of inference in order to connect the incident of not staying sober last Christmas with the incident of getting very drunk this Christmas such that *again* makes sense.

(27) a. We got some beer out of the trunk. The beer was warm.
(Haviland and Clark 1974: 514)

b. We checked the picnic supplies. The beer was warm. (ibid.: 515)

(28) a. Ed was given an alligator for his birthday. The alligator was his favorite present. (ibid.: 514)

b. Ed wanted an alligator for his birthday. The alligator was his favorite present. (ibid.: 516)

(29) Last Christmas Eugene couldn't stay sober. This Christmas he got very drunk again. (ibid.: 517)

Haviland and Clark tested this hypothesis by measuring how long participants took to understand the second sentence in cases like these. They found that comprehension indeed took longer when the antecedent was indirect, as in (27b), (28b), and (29), than when it was direct, as in (27a) and (28a), which supports the idea that extra processing is required in those cases.

Another domain where the interpretation of linguistic meaning seems to be related to inference can be seen in *coercion*. Coercion is the phenomenon by which the linguistic context surrounding some bit of an expression influences the interpretation of the meaning 'internal' to that bit of the expression, as in the following examples. *Sue slept all night* in isolation refers to one continuous instance of sleeping (30a); however, with the addition of *until she started drinking too much coffee*, it now refers to multiple instances of sleeping on separate occasions (30b). The context of (30b) can be said to "coerce" the interpretation of *Sue slept all night*. In (31), contextual coercion results in

[15] In (28a), *an alligator* refers to a specific alligator; this is what allows it to be direct. In (28b), it refers to a non-specific alligator (although an interpretation as a specific alligator is possible in theory), and thus is indirect.

the interpretation that Mary enjoyed or began reading the book (or writing it), instead of the conceivably possible interpretation that she began or enjoyed eating it, which is nonsensical in most situations (De Almeida and Dwivedi 2008: 304-305). Similarly, the interpretation of *beauty* can be coerced by context: it refers to beauty as a quality in (32a), fact-of-beauty in (32b), and is ambiguous between these in (32c).

(30) a. Sue slept all night. (Jackendoff 2002: 391)
b. Sue slept all night until she started drinking too much coffee. (ibid.)

(31) a. Mary enjoyed the book. (adapted from De Almeida and Dwivedi 2008: 304)
b. Mary began the book. (De Almeida and Dwivedi 2008: 304)

(32) a. Its beauty is sublime.
b. Its beauty was unexpected. ('The fact that it was beautiful was unexpected')
c. Its beauty is unique. ('It has a unique sort of beauty' or 'That it is beautiful is unique')

Recognizing what sort of interpretation is appropriate to a particular context involves a form of inference. Again, it depends on an understanding of how the relevant meanings function and interact.

So, the interpretation of linguistic meaning can require inference. Encyclopedic knowledge appears to be crucial for these cases of inference. Now, what we are concerned with is the instantiation of meaning construction, rather than the interpretation of meaning construction. These are not the same thing. Interpretation involves evaluating an expression in relation to linguistic and extra-linguistic context. The context can restrict what linguistic structures can be considered viable (e.g. *Its beauty is sublime* vs. *Its beauty was unexpected* or what *bank* means in *He forgot to go to the bank*), and interpretation in this sense is essentially a matter of selecting viable structures. Interpretation in terms of discursive context (like in *Ed wanted an alligator for his birthday. The alligator was his favorite present.*) is essentially a matter of identifying what is being denoted by the expression. Neither of these is quite the same thing as the instantiation of meaning construction: our problem is how expressions make sense at all. That said, the cognition underlying meaning construction does seem to be based—at least in part—on inferential processes. The navigation of paths to

meaning (as discussed in Section 3.2 above) can be achieved by inference. For example, if something is ‘here’, then it should be paid attention to; therefore ‘behold’ (c.f. (1b), (23b)); if something is not to be yielded, then it is necessary (c.f. (3c)). The problem for theoretical description becomes identifying what sort of cognition is involved in this inference.

3.3.2 Mental model cognition

My proposal regarding this—and I should stress at the outset that I am not giving a complete answer here—is that the cognition responsible for instantiating the meaning yielded by an arrangement of linguistic elements is a sort of mental model cognition. A mental model, as I mean it, is an understanding of how the elements of some system function and interact. In the case of using language for expressing meaning, knowledge of how meanings function and interact at a conceptual level serves as a background theory that informs the mechanics of meaning construction. Hampton (2017) also articulates a similar view regarding the importance of function and interaction of meanings: “the construction of complex concepts proceeds (most naturally) through the interactive combination of the intensional meanings of the individual concepts” (p. 96)

The function of this sort of mental model cognition can be easily understood through an example of a different application of it: understanding of physical mechanics and motor planning.

If I have a mug on my desk, I know that if I push it with my finger it will move, and if I push it to the edge and keep pushing it, it will fall off the desk. The understanding of the mechanics of the relevant elements constitutes the mental model. The elements that comprise the mental model are the finger, the mug, the pushing, the mug being affected by that, the moving, the edge, the falling, and perhaps the force of the pushing and feedback from that to gauge the force, friction between the mug and the table, gravity, etc. I know how they interact, so I can predict what will happen without actually doing it.

Similarly, I can throw a ball such that it arrives in an intended location. In order to do this, I must have a mental model of my biomechanics and how they relate to the

trajectory of the ball: here, the elements of the mental model are the muscular exertion of force[16], the particular technique used, and how this translates to the speed of the ball for a particular technique, the angle the ball is thrown at, the weight of the ball, the aerodynamics of the ball, etc. Again, this mental model makes the result of the throw be predictable. It should be noted that it is predictable in a flexible way: I can take into account wind or different kinds of balls (e.g. the aerodynamics of a frisbee—although this is not technically a ball— are rather different from a baseball) or different distances or heights (e.g. throwing from a lower location to a higher one), etc.

The sort of understanding I am talking about supports logical structure. For example, in a mental model of geometry, I know that if I turn left two times, I will then be oriented backward from my original facing; I know that if I turn 180 degrees once, I will be oriented backward from my original facing; and I know that if I turn 180 degrees twice, I will have the same orientation as I had originally. (Here, the elements are turns of various angles.)

In all of these cases, the key is that I know how each element functions and how it interacts with the others. This enables the result of any potential manipulation of the system to be predictable to me. In an expression, the elements are the words, morphemes, and constructions, and I know how they will interact and what the meaning of a particular arrangement will be[17]. The idea is that the cognition involved in

[16] To be precise, this also includes the exertion of *speed*: physiologically, muscles are subject to what is known as the *force-velocity relationship*, whereby higher contraction forces slow the speed of contraction, and higher contraction speed is possible only for relatively weaker contraction force. This is a main reason that throwing athletes (among others) are trained to throw with their muscles relaxed. Of course, this sort of precision is not a necessary part of a functional mental model; what is actually included in a real mental model will depend on the individual.

[17] One difference between the mental models involved in meaning construction in linguistic expressions and the motor planning examples given above is that the elements of the former are fixed to a much greater degree than are those of the latter: the elements of a particular expression are more or less determined (though there can be some amount of individual variation in analysis; c.f. Section 2.2.1.2), but for cases like the mug or ball-throwing examples, it is possible to have more or less detailed mental models with more or fewer elements.

determining what a particular arrangement means is similar to or part of the cognition involved in manipulating mental models.

The way that elements in a mental model interact open-endedly but predictably seems to be structurally analogous to the way that meanings interact with each other in linguistic expressions. Expressions can be formed from any number of arrangements of any number of elements to systematically produce effective meanings that are not fixed beforehand. This is also true of mental models: an understanding of the basic function and interaction of elements allows one to predict what will result from any arbitrary coherent arrangement of those elements. Another reason—apart from open-ended, flexible predictability—that attributing meaning construction to mental model cognition is attractive as a hypothesis is that mental model cognition supports 'reinterpretational' inference. Much of the 'reasoning' done in establishing the constructed meaning in linguistic expressions is a matter of reinterpretation, rather than of determining logical consequences. For example, deriving 'even' from 'arriving' (c.f. (23a)) or 'behold' from 'here' (c.f. (1b), (23b)) is not done on the basis that 'arriving' logically entails 'even' or 'here' logically entails 'behold', but rather by reinterpreting the meanings of 'arriving' or 'here'; in certain contexts, those meanings are equivalent to 'even' or 'behold'. This reinterpretation is based on the mechanics of the mental model. For instance, the point at which one arrives is the result and the culmination of the process of getting there, and one can arrive at points that are remarkable for one reason or another[18]; so my mother reaching the remarkable point of telling me to clean my room can be expressed as 'arriving'.

Johnson-Laird and colleagues (e.g. Johnson-Laird 1980, 1983; Johnson-Laird and Byrne 2002; Johnson-Laird and Khemlani 2014) have demonstrated that mental models can be used in inference and that positing the use of mental models can account for what sorts of problems people tend to find difficult and what sorts of problems they tend to find easy. Now, it must be pointed out that what I mean by 'mental model' differs slightly but importantly from what Johnson-Laird and colleagues mean by the term. Johnson-Laird's conception has mental models as iconic abstract dioramas of

[18] This information is (some of) the encyclopedic knowledge that constitutes the mental model.

states of affairs. For example, the meaning of *Some of the actors are bakers* can be represented as a mental model (in his sense) to the effect of:

Actor	Baker
Actor	Baker
Actor	
	Baker

(Johnson-Laird and Khemlani 2014: 34)

Manipulating this sort of diorama can be used to make inferences. Thus, an inference like *A is to the left of B; C is to the right of B: therefore A is to the left of C* is made quite easily; the diorama is quite simple:

A B C

People tend to make certain kinds of errors in reasoning. For example, to the problem

> Al is a blood relative of Ben.
> Ben is a blood relative of Cath.
> Is Al a blood relative of Cath?,

people tend to respond *yes* (Johnson-Laird 2010: 18244); now, this is not actually a valid inference, since Al and Cath could be Ben's parents, and thus blood relatives of Ben, but not of each other. Johnson-Laird remarks that it seems that this error is produced because people tend to make a single model with properties like those of systems of lineal descendants or filial relations, in which it *would* be true that Al is a blood relative of Cath.

Similarly, given the problem

> All of the Frenchmen in the restaurant are gourmets.
> Some of the gourmets in the restaurant are wine-drinkers.
> What, if anything, follows?,

participants tended to respond

Therefore, some of the Frenchmen in the restaurant are wine-drinkers.

This inference is not valid (it is possible for there to be non-French gourmets, and therefore for all of the wine-drinkers to be non-French), but it is well-accounted for by Johnson-Laird's theory. The diorama that the participants construct is something of the following form:

Frenchman	gourmet	wine-drinker
Frenchman	gourmet	wine-drinker
Frenchman	gourmet	

This is constructed by overlaying structures for the two premises (all the Frenchmen are gourmets and some of the gourmets are wine-drinkers). This does yield the conclusion that some of the Frenchmen are wine-drinkers. However, for an analogous problem like

All of the Frenchmen in the restaurant are gourmets.
Some of the gourmets in the restaurant are Italians.
What, if anything, follows?,

participants did not tend to respond

Some of the Frenchmen are Italians.

In their mental model, a person cannot simultaneously be a Frenchman and an Italian, and so the corresponding diorama

Frenchman	gourmet	Italian
Frenchman	gourmet	Italian
Frenchman	gourmet	

does not work; instead, an alternative diorama consistent with the premises must be constructed:

Frenchman	gourmet
Frenchman	gourmet
Frenchman	gourmet
gourmet	Italian
gourmet	Italian

Here, no Frenchman is an Italian, and the premises that all Frenchmen are gourmets and some of the gourmets are Italians are captured (Johnson-Laird 2010: 18245-18246).

Now, my *mental model* is more abstract: it refers to an understanding of how the elements in a system function and interact, and not to a diorama. Whereas Johnson-Laird envisions a mental model as a particular configuration or state, a mental model in my sense is the manipulable system itself. It should be noted that my definition is not incompatible with Johnson-Laird's sense—since his sense corresponds to the mental representation of the result of a particular manipulation, it can be considered to be subsumed in mine as a special case. With my definition, I do not need to assume a particular form of mental representation, which is better for generality. I am unwilling to assume that the meanings of linguistic expressions are mentally represented in the sort of iconic form Johnson-Laird describes, but I will not rule out that they could be in some cases and for some individuals. Also, my definition focuses on the crucial mechanism: the understanding of function an interaction of the elements is what forms the basis for any manipulation of a mental model—in my sense or in Johnson-Laird's (and it would be required for setting up any diorama)—and it is what gives it its predictability; this point will be discussed more in the next subsection.

3.3.3 Simulation

A similar proposal to mine is that meaning in linguistic expressions is created as simulations (e.g. Barsalou et al. 2008; Evans 2009, 2016). These theories are predicated on a perceptual format of the mental representation of concepts (c.f. Section 2.2.2 (Modal versus amodal representation)). Essentially, concepts are mentally represented through the activation of perceptual representations of sensorimotor and introspective

experience. These perceptual representations can then be activated in the absence of a perceptual stimulus. This re-activation of perceptual representations is called *simulation*. The idea for simulation semantics is that the meaning of linguistic expressions is mentally represented as simulated experience, so the mental representation of the meaning of a particular expression is a constructed perceptual image of it (though not necessarily a conscious image). For example, the meaning *The actress put on her red lipstick* would be mentally represented by a perceptual simulation of that scenario (Evans 2016: 5). Evans notes that the particular shade of red included in such a simulation would be different from that in, e.g., *The red fox jumped over the stream* (ibid.). The word *red* by itself does not indicate the precise hue; rather, encyclopedic knowledge (here, pertaining to the color of red lipstick and red foxes) guides the simulation.

Now, there are actually two hypotheses underlying the view that meaning is constructed through simulation: a hypothesis about the format of the mental representation of meaning, and a hypothesis about how meaning is constructed.

There are two related aspects that comprise the format hypothesis: that linguistic concepts are represented perceptually, and that meaning in expressions is created as simulated experience. I am unwilling to assume that either of these is correct. As I discussed in Section 2.2.2, there are two major problems for hypothesizing that concepts are represented through perceptual information. One is that perceptual information (alone) may be unable to instantiate concepts. Abstract concepts are particularly problematic, but the mental representation of even concrete concepts may not be instantiated through perceptual information. The other is that it is unclear whether the numerous experimental results relating to this issue support the hypothesis. It is not clear whether the sensorimotor neural activity found experimentally to occur in the context of conceptual processing is involved in constituting the mental representation of concepts or is secondary, and associatively triggered or indicative of processing the concepts rather than representing them. Also, it is uncertain whether the brain region in which a particular instance of neural activity occurs should be taken to be indicative of mental representation in a particular format. Furthermore, even if sensorimotor representations are involved in building the mental representations of

concepts, their connection to the conceptual representation may be quite indirect, such that the conceptual representation cannot really be considered to be perceptual in format. Pure sensory information is certainly inadequate for constituting conceptual representations[19]; at a minimum, the sensory information would have to be processed or interpreted. At these levels where processed sensory information is represented, the mental representation is more abstract, and at some point, the format of representation cannot be considered perceptual. I am also skeptical about the idea that the meaning of expressions is represented as simulated experience: I am unconvinced that simulation necessarily occurs when mentally representing the meaning of expressions, and even if it does occur, it may be secondary. Of course, the same problems that exist for the perceptual representation of concepts noted above apply here as well; it is doubtful whether a simulation can actually carry the meaning of an expression[20].

With regard to the hypothesis about how meaning is constructed cognitively, simulation semantics is similar to my proposal described in Section 3.3.2 that the cognition involved in making sense of mental models is related to that involved in making sense of arrangements of linguistic elements. The same cognitive mechanisms—determining meaning in mental models from knowledge of the function and interaction of the elements—presumably would be involved in constructing simulations: in order to know what sort of simulation to construct, one must know the meaning. But still it misses the point we are concerned with. Essentially, simulations do not establish what meaning is constructed; rather, they are dependent on the meaning already being established. The process of creating a simulation necessarily depends on the mechanisms that I propose. The following analogy illustrates this point: imagine that someone is considering a system consisting of two linked gears, and wants to figure out what direction one of the gears will rotate in if the other is turned; they might make a mental simulation of the rotation of the gears, and then determine the direction of rotation by observing this simulation. But in order to set up the simulation correctly, they must already understand the mechanics that actually govern how the gear will

[19] With the exception of perceptual concepts, described in Sections 2.2.2 and 2.2.3.2.

[20] That is: even if a simulation *corresponds* to the meaning of an expression, it probably does not *constitute* its mental representation.

rotate (i.e. the physics and geometry of how rotation functions and how objects interact when they move into each other). It is this understanding that is actually responsible for cognitively determining the direction of rotation. That is the reason that I emphasize the understanding of function and interaction as the crucial mechanism. Hampton (2017: 112-13) reports a study where participants had to describe concepts such as 'a bird that is also a kitchen utensil'; he notes that certain emergent properties (like that a 'woodpecker whisk' would not need electricity and therefore would be convenient for camping, but may not be hygienic) arise here, and attributes this to simulation. But really they come out of the understanding of function and interaction that I emphasize; the mental process of simulation is not required to produce them. If a simulation is made, it is based on the mechanisms that I emphasize. Now, I am not at all arguing against the use of simulation in cognition. Rather, I am shifting the emphasis to what has to be the fundamental mechanism. The predictability of what meaning will be yielded by a particular arrangement of linguistic elements comes from the understanding of the mechanics of meaning construction.

3.4 Summary

Linguistic expressions are not looked up; rather, they must be created by arranging linguistic elements. In order for this sort of system to be capable of expressing thought, the meanings yielded by particular arrangements of elements must be predictable to the language user.

In Section 3.1, we discussed apparent degeneracy in this predictability. There are two basic cases of this. One is caused by mismatches between the formal structure of an expression and what are actually functioning as meaning-bearing linguistic elements. The other relates to underdetermination of meaning; here, if the meaning expressed by a particular linguistic form is one of several distinguished meanings corresponding to that form, there are in fact several linguistic structures that are distinct from each other, and if it is not, then the meaning of the arrangement is simply schematic. We saw that neither of these cases is problematic for predictability.

In Section 3.2, we looked at the dynamicity of meaning construction; the meaning yielded by a particular arrangement is not just a structured list of the component elements. Particular expressions take particular 'paths' to the meaning that they express, and approximately-equal meanings can be obtained from wildly different paths. Differences in the meaning of linguistic elements across languages mean that often, closely-corresponding elements are used similarly in some situations, but differently in others. Another aspect of this dynamicity can be seen in the fact that approximately-equal meanings can often be expressed through either atomic linguistic representations or composite linguistic representations; this reflects some of the fluidity in the way meaning is expressed in language. Additionally, meaning construction is important in combinatorial arrangements of multiple elements, but also can be involved in deriving meaning from single-element arrangements.

In Section 3.3, we considered the question of what cognitive mechanisms are responsible for establishing the meaning yielded by an arrangement. We first noted that meaning construction appears to be related to inference, although inference for the interpretation of meaning is distinct from inference for the establishment of meaning. We then described a proposal suggesting that the establishment of meaning is achieved through a sort of mental model cognition, and distinguished this proposal from the mental model theory of Johnson-Laird and theories of simulation semantics. The basis of meaning construction is the understanding of how linguistic meanings function and interact; application of this understanding is probably closely related to the cognition involved in making sense of mental models and to that involved in constructing simulations. I am under no illusion that this can be considered a complete answer—it is still quite vague, and the details of what specific cognitive operations are involved in its implementation need to be worked out. But this characterization is essentially important. Its fundamental significance is that the understanding of the mechanics of meaning construction can function as a background theory on the basis of which the meaning yielded by a particular arrangement is predictable. This is crucial for the remaining parts of the theory, laid out in Chapters 4 and 5.

Chapter 4

Constructing an Arrangement of Linguistic Elements

The next part of the problem we will consider is the issue of what constitutes the task of constructing an arrangement of linguistic elements to express a thought. A linguistic expression is made up of a set of elements that are structured in some particular fashion. In the context of deriving a linguistic expression from a language-independent thought, the conceptual material in the thought must be related to this arrangement. Since the original thought is language-independent, it does not have the particular properties of the arrangement, so the particular elements of the expression and how they are structured are undecided initially. Composing an arrangement entails deciding the elements and their structure. This is done by taking the conceptual material in the thought and organizing it into the particular elements and the structure they have in the arrangement. The circularity of this description is important: determining what the elements are and how they are structured are interrelated problems. What I refer to as 'organization' applies to the conceptual material that is to be expressed; this is distinguishable from the actual selection of elements. But really these aspects are not separable, since the language user has to organize the conceptual material into the actual elements that they use in the expression. Fundamentally important for this organization of the conceptual material in the thought is the understanding of the mechanics of meaning construction. As articulated in Chapter 3, the language user's theory of meaning construction—of how individual meanings function and interact—allows the meaning yielded by a particular arrangement to be predictable to them, and thus allows them to suitably organize the material. I call the organization of conceptual material in accordance with the mechanisms of meaning construction *compartmentalization*. This is a key process for the expression of thought in language.

In order to create an arrangement that expresses their thought, the language user must compartmentalize the thought.

In Section 4.1, I discuss the organization of conceptual material. First, I review theories on the organization of semantic information in expressions. These describe some important aspects of the issue we are concerned with here. Following this, I introduce the theory of compartmentalization. To give an understanding of how compartmentalization works, I describe basic patterns of variation in how thoughts can be compartmentalized. I also show that because of the basic structure of linguistic expressions, there are often interdependencies between different elements in an arrangement; compartmentalization must therefore often take into account multiple bits of conceptual material at once.

In Section 4.2, I then discuss the issue of the selection of linguistic elements. The typical explanation is that this is done through the process of categorization; I discuss the details of categorization and how it is more complex than one might assume. Secondly, I summarize psycholinguistic research on mentally accessing elements. This research has established that mental representations of linguistic elements form a network with connections of various strengths between its members; the cognitive mechanisms of the access of elements are associative to a significant degree. Finally, I argue that while categorization is important for some aspects of the process of selection of elements, it does not account for others, and that the process of selection of elements should be understood through compartmentalization.

In Section 4.3, I give a synopsis of my account of this aspect of the problem of expressing thought in language.

4.1 Organization of information in expressions

4.1.1 Theories on the organization of semantic information

In this section, I describe three significant existing theoretical accounts of the organization of semantic/conceptual information in linguistic expression.

4.1.1.1 Grouping

Langacker (1997, 2000, 2008, 2009b) proposes the notion of *grouping* to describe the conceptual process by which separate cognitive objects function as a single unit.

Different patterns of conceptual grouping can be depicted in linguistic relationships. Grouping is one aspect of the structural relationships that can be expressed in language (c.f. Section 2.2.1.4.4).

For example, it is possible to conceptualize *drinking and smoking* as separate objects or as one object. In (1), which of these conceptualizations is depicted is revealed by number agreement in the verb: treatment as a plural corresponds to a non-grouped conceptualization (1a), and treatment as a singular corresponds to a grouped conceptualization (1b).

(1) a. [Drinking] and [smoking] do not improve your health. (2009b: 52)
b. [Drinking and smoking] does not improve your health. (ibid.)

The same is true for *a flock of geese*, which can refer to the individual geese that form the flock (2a) or to the flock as a whole (2b).

(2) a. A flock of geese were flying overhead. (ibid.)
b. A flock of geese was flying overhead. (ibid.)

While the conceptual objects that may be grouped or ungrouped can correspond to distinct linguistic elements, as in (1) (such that I can indicate the grouping with bracketing), this is not always the case: in (2), the question is not whether [flock] is grouped with [geese], but whether the geese are distinguished as individuals or not. The workings of this sort of conceptual grouping is further illustrated by the fact that if the above are combined with material that requires conceptualizing individual geese, while the expression corresponding to the conceptualization of separate geese (3a) is perfectly fine, that corresponding to the conceptualization of the undivided flock (3b) is incongruous.

(3) a. A flock of geese are flying overhead by flapping their wings. (2009b: 363)
b. *A flock of geese is flying overhead by flapping {its / their} wings. (ibid.)

Langacker also notes that in other linguistic contexts (e.g. *Drinking and smoking will not improve your health* or *We saw a flock of geese flying overhead*), the above distinctions in conceptual grouping are not reflected in the linguistic form (2009b: 52). Likewise, *Alice and Bill resigned* can depict either two separate events (Alice resigning and Bill resigning) or a single event (Alice and Bill resigning together) (2009b: 359-362). *The boxes are heavy* does not distinguish between whether the boxes are heavy individually or collectively (2009b: 360 note 6; also see p. 51). Langacker remarks that in these cases, the language user may not make a distinction in their mental representation of the meaning of the expression either (i.e. in their analysis, the meaning can be relatively vague in this respect) (2009b: 361).

But in some cases, the grouping is fixed. For instance, *gin and tonic* (as an established expression) does not refer to [gin] and [tonic] as separate objects; *I drank gin and tonic* does not mean 'I drank gin and I drank tonic' (2009b: 365). Whereas a *red shirt* is generally red as a whole, a *red and yellow shirt* is not red as whole and also yellow as a whole (2009b: 363).

The conceptual organization related to grouping can also involve the attachment of some (collections of) cognitive objects to others. A contrast can be seen in the following example. In (4a), *intelligent* is a property possessed separately by each of *those women* as individuals. This is diagrammed on the left half of Figure 7: the circles represent the women, and the squares represent the property 'intelligent'; the attachment is to each woman individually; the rounded rectangle indicates that the women still form a group, referable to as *those women* (e.g. it is not *That woman is intelligent, and that woman is intelligent...*). In (4b), however, *numerous* is a property of *the problems* as a whole. This is diagrammed on the right half of Figure 7: the circles represent the problems, and the square represents the property 'numerous'; the attachment is to the collection of problems as a whole (indicated by the rounded rectangle). The different manners of attachment of *intelligent* and *numerous* are part of the differences in the organization of the conceptual material depicted in the expressions.

(4) a. Those women are intelligent. (Langacker 2009b: 50)

b. The problems with that idea are numerous. (ibid.)

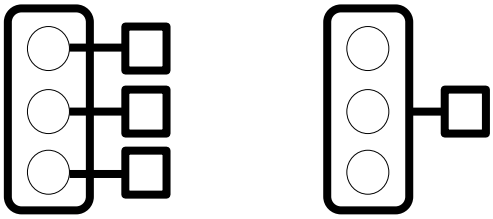

Figure 7 (adapted from Langacker 2009b: 50)

Grouping can also be made on the basis of focus or contrast. For example, in (5), there are two 'dimensions' of contrast. One is *Jack* versus *Jill* (bolded below), and the other is *whisky* versus *gin* (underlined below). Langacker analyzes the elements that would be pronounced with stress (here, *Jill* and *gin*) as forming a grouping (also see Langacker 1997: 15).

(5) **Jack** likes <u>whisky</u>. **Jill** prefers <u>gin</u>. (2009b: 348)

In terms of structural relationships in the mental representation of the meaning of the construction (if contrast is indeed intended), there also is an association between *Jack* and *whisky* to parallel that between *Jill* and *gin* (the only grouping that Langacker mentions is that related to prosodic stress), and there are additionally associations between *Jack* and *Jill* and between *whisky* and *gin*. I do not know if Langacker himself would consider these to be groupings, but the conceptual organization involved does appear to be quite similar.

Langacker also points out that grouped elements do not need to contiguous; they can be syntactically disparate, as in (6) below.

(6) The <u>headway</u> that we managed to <u>make</u>… (Langacker 2000: 157)

The same applies to the grouping in (5) above. Relatedly, the differences in conceptual grouping in (4) are not reflected in differences in syntax. The grouping in the conceptualization depicted in an expression does not need to match the syntactic form of the expression.

Additionally, conceptual material organized into mental spaces can be considered to be grouped together. In example (7), there are two mental spaces: the space of reality (bolded), and the space of Jan's belief (underlined). The reality space contains the conceptual material relating to Jan's son's actual degree of intelligence, and the belief space contains the conceptual material relating to what Jan believes is his degree of intelligence. The conceptual material belonging to each space can be said to be grouped together (in this case, the linguistic structure that the two spaces correspond to overlap; *smarter* is simultaneously grouped into both).

(7) Jan believes that <u>her son is **smarter**</u> **than he really is**. (Langacker 1997: 15; also discussed earlier in this dissertation in Section 2.2.1.4.4)

4.1.1.2 Individuation and propositionalization

Chafe (1975, 1977, 2005) and Croft (2007) develop a model of verbalization. Chafe (1975, 1977) approaches the question of how experience can be verbalized from the standpoint of discourse analysis.

He proposes that first, the whole experience that is to be expressed (e.g. the content of a story, etc.) is divided into scenes according to certain prototypical patterns or "schemas". For example, in one particular schema, "protagonist A is engaged in some background activity as the episode opens. Then comes protagonist B's arrival on the scene. There follows a conversation between the two. Then some action is taken, usually by protagonist A but affecting protagonist B, and finally there is the departure of the latter." (1975: 90). These schemas are used repeatedly in many stories. He calls this initial organization *subchunking* or *schematizing*.

When subchunking has narrowed the focus down to an "event" (basically, a chunk that can be expressed in a sentence), the next task, as identified by Chafe, is to

propositionalize the scene (this is also referred to as *framing*) by identifying the participants of the event and their roles.[1]

In order for a proposition to be determined, the elements that participate in that proposition must also be determined. Accordingly, when some event is being construed for the purposes of verbalization, it must be divided into the participant elements. This is called *individuation*. As shown in example (8), different individuations are possible for a given event, depending on the construal. In the construal depicted in (8a), *he* and *a girl* are identified as (separate) participants; in that of (8b), instead, *both of them* is identified as a (single) participant.

(8) a. he passes a girl on a bicycle (Croft 2007: 356)
b. you see both of them converging (ibid.)

In the final step, the objects thus individuated are mapped to lexical items via categorization (c.f. Section 4.2.1 below), which completes the verbalization process.

Additionally, a goal for Chafe (2005) and Croft (2007) is to provide a rationalization for the existence of grammar in terms of conceptual processes. To the basic model described above, Croft adds the processes of *particularizing*, which grounds (c.f. footnote 43 in Section 2.2.1.4.2) the individuated elements, and those of *structuring* and *cohering*, which tie those elements together using grammatical content[2].

With regard to the first of these, *particularization*, the content to be expressed often corresponds to "natural" individuations[3]. But in other cases, the content to be

[1] Positing propositionalization as a stage pre-supposes that the conceptual content to be expressed is originally not mentally represented in a propositional form (or at least that it does not need to be) (recall from Section 2.2.1.4.1 that Chafe and Croft consider propositions to be construal-dependent). I am agnostic on what the format of the original thought to be expressed is, but I do assume that whatever its form, the thought is generally not the same as what is ultimately expressed in language. See Section 2.1.

[2] I am using Croft's terminology here. Chafe's *orientation* is similar to Croft's *particularizing* and *cohering*, and Chafe's *combination* is similar to Croft's *structuring*.

[3] Croft does not give an explicit definition of what he means by "natural", but it appears that a "natural individuation" refers to conceptual content that can be expressed purely through lexical elements without needing any grammatical elements.

expressed does not match a natural individuation; in these cases, grammatical elements such as numerals or quantifiers and number (as in (9a)) or expressions of aspect (as in (9b)) must be added (Croft 2007: 358-359). Croft refers to this as *selecting*.

(9) a. six birds (Croft 2007: 358)
b. I started to eat (ibid.: 359)

Also, another part of particularizing involves *situating* the conceptual content using deictic elements, such as demonstratives and tense expressions, and modals (ibid.: 360-361)[4]. Examples are shown below.

(10) a. that chair (ibid.: 360)
b. She ate (ibid.)
c. The bird might be in the tree (ibid.: 361)

The other processes, *structuring* and *cohering*, are responsible for linguistically encoding aspects that relate to constructional meaning. *Structuring* concerns determining parts of speech, argument structure, and pragmatic information structure[5] (ibid.: 363-373). Croft points out that this is language-specific: for example, English distinguishes between adjectives and verbs (e.g. *She is tall* vs. *She is a student*), but other languages do not (p. 364). *Cohering* concerns links between clauses and reference tracking (ibid.: 373-376). Clause linkage is achieved by, for example, the underlined elements in *So they're walking along, and they brush off their pears, and they start eating it* (p. 373); reference tracking is achieved, for example, in the constructions used in *Sally peeled and ate the banana*, where the subject and object are the same for both actions, and *Sally ate the banana and Gary the watermelon*, where both the subject and object are different across the actions (p. 374).

[4] Croft (ibid.) also identifies situating with organization into mental spaces.

[5] Pragmatic information structure refers to what is "given information" and "new information" in a discourse. This is not to be confused with what I refer to in this dissertation as the 'organization of information', which concerns the conceptual material in the thought and in the mental representation of the linguistic expression.

With regard to my own theory, individuation is an important aspect of compartmentalization, but I do not assume that the various conceptual processes Croft posits beyond that (particularizing, structuring, and cohering) actually have distinct mechanisms beyond differences in the functioning of the particular meanings of the elements involved. Also, subchunking has to do with discourse planning, an issue outside the scope of this dissertation; as discussed in Section 2.4 (also see the introduction to Section 2.1), our problem starts at the point whenever the thought to be expressed is formed.

4.1.1.3 Conflation

Talmy (2000b) discusses phenomena like those below that concern what semantic aspects are encoded in particular linguistic elements and, cross-linguistically, whether certain semantic aspects are encoded in one element or across several elements. He refers to the phenomenon of multiple aspects being simultaneously encoded in a single element as *conflation*.

Talmy distinguishes between (fact of) movement and manner of movement. In examples (11) and (12) below, English *floated* encodes both. In Spanish, however, movement and manner are not encoded in the same element; they must be separated out into *entró* and *flotando* (11b) or *salió* and *flotando* (12b). It should be noted that *flotando* ('floating') can be omitted (with the resulting meaning being *the bottle entered/exited the cave*), but *float* cannot.

(11) a. The bottle floated into the cave. (Talmy 2000b: 49)
b. La botella entró a la cueva (flotando). (ibid.)
the bottle moved.in to the cave (floating) (ibid.)

(12) a. The bottle floated out of the cave. (ibid.)
b. La botella salió de la cueva (flotando). (ibid.)
the bottle moved.out from the cave (floating) (ibid.)

Spanish consistently separates manner from movement where English can conflate them into one word. Conversely, Spanish often conflates motion with path,

where English separates them. This can be seen in (13) and (14) below. In (13), *rolled* can include information about moving the keg and how it was moved, whereas these must be expressed in separate words (*metí* and *rodándolo*) in Spanish. However, *rolled* has no information about where the barrel went; by contrast, *metí* does: it was put in. Likewise, in (14) *twisted* can express that I removed the cork and simultaneously how I removed it. In Spanish, these are expressed separately, by *saqué* and *retorciéndolo*. If one wishes to say something like 'twist the cork', the twisting and the removing must still be separated, as in (14c). Similarly to *metí*, *saqué* includes information about where the cork went (i.e. out), where *twisted* does not (one can *twist the cork into the bottle* or *twist the cork out of the bottle*; the distinction does not lie in *twist*).

(13) a. I rolled the keg into the storeroom. (ibid.: 51)
b. Metí el barril a la bodega rodándolo (ibid.)
I.put.in the keg to the storeroom rolling.it

(14) a. I twisted the cork out of the bottle. (ibid.)
b. Saqué el corcho de la botella retorciéndolo. (ibid.)
I.took.out the cork from the bottle twisting.it
c. Retorcí el corcho y lo saqué de la botella. (ibid.)
I.twisted the cork and it I.took.out from the bottle

Information about the path and destination[6] of movement can be conflated as well. For example, Atsugewi has many terms that would correspond to *into X* in English, with each term specifying a different X. The examples in (15) are from Talmy (2000b: 110-111). Some of them must be followed by *-im* or *-ik·*, which Talmy glosses as 'hither' and 'thither', respectively; these are indicated by a +.

(15) a. -iċt 'into a liquid'
b. -cis 'into a fire'
c. -isp-u· + 'into an aggregate' (e.g. bushes/a crowd, a rib cage)

[6] Talmy uses the term *ground*, which means essentially 'reference object'. While this terminology has the advantage of being more generalized (it is able to include origin of movement, for instance), I use 'destination' here because it is clear without an explanation.

d. -wam ‘down into a gravitic container’ (e.g. a basket/cupped hand/pocket)
e. -wamm ‘into an areal enclosure’ (e.g. a corral/field/area occupied by a pool of water)
f. -ipsn^{u} + ‘(horizontally) into a volume enclosure’ (e.g. a house/oven/crevice)
g. -tip-u· + ‘down into a (large) volume enclosure in the ground’ (e.g. a cellar)
h. -ikn + ‘over the rim into a volume enclosure’ (e.g. a gopher hole/mouth)
i. -ikc ‘into a passageway so as to cause blockage’ (e.g. choking/shutting/walling off)
j. -iḱs^{u} + ‘into a corner’ (e.g. a room corner/edge between wall and floor)
h. -mik· ‘into the face/eye (or onto the head) of someone’
i. -miċ ‘down into (or onto) the ground’
k. -cis^{u} + ‘down into (or onto) an object above the ground’ (e.g. the top of a tree stump)
l. -iḱs ‘horizontally into (or onto) an object above the ground’ (e.g. the side of a tree stump)

Another sort of conflation is that of motion and the nature of the object involved. Atsuwegi uses different verb roots for expressing the motion of different kinds of objects. The examples in (16) below are from Talmy (2000b: 58).

(16) a. -lup- ‘for a small shiny spherical object (e.g. a round candy) to move/be located’
b. -ṫ- ‘for a smallish planar object that can be functionally affixed (e.g. a stamp) to move/be located’
c. -caq- ‘for a slimy lumpish object (e.g. a toad) to move/be located’
d. -swal- ‘for a limp linear object suspended by one end (e.g. a shirt on a clothesline) to move/be located’
e. -qput- ‘for loose dry dirt to move/be located’
f. -sṫaq̇- ‘for runny icky material (e.g. mud) to move/be located’

Talmy also remarks (2000a: 95-96 note 18) that one can use the (somewhat tortured) expression *a clustering of a set of trees*; this expression separates the three

aspects of what kind of thing is in question (*trees*), the fact that it is a bit of that thing (*a set*), and the shape of that bit (*a clustering*). It is also possible to use the single element *grove*, which conflates the kind of thing (trees), the fact that it is some bit of that thing, and the shape of that bit (a "clustering form"). The same applies to *drop*, which in his analysis conflates the kind of thing (liquid), the fact that it is some bit of that thing, and the shape of that bit (a "small globular form") (c.f. *a small globular form of an amount of liquid*, where they are separated)[7].

Nuyts (2012) describes an analysis of expressions of epistemic modality like those in (17) below. In his analysis, these expressions encode an epistemic evaluation simultaneously with other aspects; a difference between these three expressions is that (17a) expresses the (basis for the) evaluation of likelihood as the speaker's subjective judgment, while (17b) expresses it as an intersubjective judgment (i.e. one "shared by a wider group of people, possibly including the hearer" (p. 326)), and (17c) "conveys neither subjectivity nor objectivity" (ibid.). This kind of analysis also can be considered essentially a conflation.

(17) a. I <u>believe</u> he went to the bakery (Nuyts 2012: 326)
b. It is quite <u>likely</u> that he went to the bakery (ibid.)
c. He <u>probably</u> went to the bakery (ibid.)

In an important experiment, Kita and Özyürek (2003) showed that conflation in linguistic expressions can be correlated with conflation in accompanying spontaneous gestures. They examined linguistic and gestural encoding of descriptions of swinging and rolling scenes (participants explained scenes in a cartoon) in English, Japanese, and Turkish speakers.

Japanese and Turkish have no term for 'swing' (and no "readily accessible paraphrase" either (p. 13)), and so the linguistic encoding ('go' / 'fly' / 'jump') "omitted" information related to the shape of the trajectory that is present in English

[7] Actually, Talmy says that *drop* conflates only the second two (the bit of the thing and the shape of the bit), and *tear*, by contrast, conflates all three (including the kind of thing). I fail to see the distinction: while *tear* does indeed specify the particular kind of liquid, liquid itself (specified by *drop*) surely counts as a kind of thing.

expressions like *swung across the street*; English speakers almost exclusively used arc gestures; Japanese and Turkish speakers used a mix of arc gestures, arc and straight gestures, and straight gestures.[8]

With regard to 'roll', the organizational structure of the English *rolls down the street* and its Japanese and Turkish equivalents (*korogatte saka o kudaru*[9] / *yuvarlan-arak cadde-den iniyor*) (lit. 'descends the street rolling') are different. English conflates manner and movement together (*rolls*) and separates this from path (*down*) (though, as will be discussed in Section 4.1.2.1, this can depend on the individual's analysis), whereas Turkish and Japanese conflate information about movement and path into one element ('descend') and separate this from information about manner ('rolling'). English speakers tended to use manner/path conflated gestures, whereas Japanese and Turkish speakers tended to use manner-only and path-only gestures. Gestural representation necessarily expresses movement and path simultaneously, so this gestural encoding follows the organizational structure of the linguistic encoding.

4.1.2 Compartmentalization

The above theories describe some important aspects of the organization of information in linguistic expressions. These aspects, however, only account for part of the issue of how conceptual material in language-independent thought is organized for expression in language. In order to express a thought, this conceptual material must be organized into the structure of the expression—into an arrangement composed of available linguistic elements that yields a suitable meaning. This is the notion of compartmentalization (Polo-Sherk 2020). Compartmentalization is a (re-) conceptualization process: compartmentalizing a thought amounts to a re-conceptualization of it. One obvious constraint is that the thought must be re-

[8] Strictly speaking, the issue of the omission or inclusion of information in an expression is not a matter of conflation (which is restricted to whether a particular element encodes multiple semantic aspects). This issue is, however, included in the domain of compartmentalization, introduced in Section 4.1.2 below.

[9] This actually means 'roll down the hill' (rather than 'the street').

conceptualized in terms of meanings associated with elements that exist in the target language. Another is that the arrangement of those meanings yields an approximation of the original thought.[10] Meaning construction is central to compartmentalization: the units into which the thought is divided and their organization are determined in large part by the language user's understanding of the mechanics of meaning construction (c.f. Chapter 3); that is to say that knowing what meaning will be yielded by a particular arrangement guides the compartmentalization of the thought into that arrangement.

This subsection has two parts. In 4.1.2.1, I discuss basic patterns of variation in compartmentalization. I have two goals here. One is to show how compartmentalization works by discussing different ways information can be compartmentalized. The other is to impart an awareness of the great variety in how a particular piece of material can be compartmentalized that exists. The way a particular piece of material is compartmentalized is influenced by what is required and permitted by different languages and in different linguistic contexts. This last point is also a theme of Section 4.2.1.2. In this second part, I explain that different parts of an expression interact with each other in constructing meaning, and so compartmentalization must be made in consideration of multiple bits of conceptual content. I discuss a large number of examples here, but they all serve to illustrate these basic points.

4.1.2.1 Basic patterns of variation in compartmental structure

Here, I describe prototypical patterns in variation of the compartmentalization of conceptual material. I include four such patterns here: *relative shifting*, *relative separation/compression*, *indexation vs. word position*, and *fundamental reformulation*.

As I describe them here, these patterns exist between alternative compartmentalizations in different expressions (they are relative variations in this sense). These are not, strictly, 'ways to compartmentalize a thought'. Presumably,

[10] One may identify a further (redundant) constraint: that the elements be compatible with each other. If certain elements are incompatible with each other, then this reduces which elements from the target language are available; at the same time, incompatible elements cannot yield the correct meaning.

however, operations to these effects can occur during the compartmentalization of the thought (i.e. a cognitive object in the thought can be shifted or separated; several cognitive objects can be compressed; the thought or parts of it can be fundamentally reformulated)[11].

I refer to these patterns as 'prototypical' because different compartmentalizations often exhibit mixtures of these patterns (e.g. both shifting and separation/compression), and because fundamental reformulation in particular is a matter of degree.

4.1.2.1.1 Relative Shifting

Relative shifting concerns differences in the structural position that elements corresponding to a particular bit of information are placed in.

The English and Japanese expressions in (18), which have approximately equal meanings, negate different parts of the expression. In the English (18a), the person who may be here or not (i.e. the subject) is negated; in the Japanese (18b), by contrast, whether they are here or not (i.e. the verb) is negated. (Also, note the path to meaning involved in generating 'anyone' from 'who' + ADDITIVE (c.f. example (50) below). Here, the negation is shifted.

(18) a. No one is here.
b. koko ni dare mo inai.
here LOC who ADDITIVE be.NEG
(lit. 'anyone is not here')

Relative shifting also occurs between different styles of possession marking. Some languages mark the possessor, and others the possessee, as shown in the

[11] Indexation vs. word position does not have the generality of the other patterns; were I to include it in this list as reflecting an operation on the thought, I would have to assume that the language-independent thought is structured in such a way that it has analogs to word position and indexation, and (as discussed in Section 2.1) I am unwilling to make such assumptions. For the other patterns, I only need to assume (at most) the potential existence of 'cognitive objects' in the original thought, which I do think is reasonable.

examples below (Dryer 2007: 178-179). The New Guinean language Hua (19a) marks the possessor; the Canadian language Haida (19b) marks the possessee.

(19) a. de-ma' fu
man-GEN pig
'the man's pig' (ibid., citing Haiman 1980)
b. Wā´nəgən gi´t-g̣a
Wanagan son-POSSESSED
'Wanagan's son' (Dryer 2007: 179, citing Swanton 1911)

Another example is seen in variation between adjectival (i.e. to the noun) and adverbial (i.e. to the verb) attachment of bits of semantic material in certain cases. In cases like (20), the only conceptualization that can be depicted in English has the tastiness attached to the noun (20a) (c.f. **Bake cookies tastily*); in Japanese, however, it is also possible to compartmentalize the relevant conceptual material such that the tastiness is attached to the verb (20b); this has more or less the same meaning as the English in (20a).

(20) a. Bake tasty cookies
b. kukkii o oishi-ku yaku
cookie ACC tasty-ADV bake
(lit. 'bake cookies tastily')
c. oishi-i kukkii o yaku
tasty-ADJ cookie ACC bake
(lit. 'bake tasty cookies')

Another sort of shifting is that between overt elements and non-overt constructions. Whereas, for example, the English construction [A's B] and the examples in (21) from Japanese (21a) and Hebrew (21b) use an overt element, in many other languages, possession is compartmentalized purely into a syntactic arrangement of the form [NOUN1 NOUN2], with no morphology explicitly indicating possession. The Niger-Congoan language Yoruba (22a), the New Guinean language Kobon (22b), and

Hebrew (22c) all use this style[12]. The shifting here is between encoding in an overt element and encoding in the syntactic arrangement: while in the examples described above in (18)-(20), all three overt elements (e.g. [man] ['s] [pig]) are preserved between alternative compartmentalizations, in the following cases, the overt element (['s]) exists in one alternative and does not exist in the other[13].

(21) a. Taro no kuruma
Taro GEN car
'Taro's car'
b. ha-dira shel moshe
the-apartment of Moshe
'Moshe's apartment' (Glinert 1989: 34)

(22) a. fílá Àkàndé (lit. 'cap Akande'; 'Akande's cap') (Croft 2003: 32, citing Rowlands 1969: 44)
b. Dumnab ram (lit. 'Dumnab house'; 'Dumnab's house) (ibid., citing Davies 1981: 57)
c. dirat Moshe (lit. 'apartment Moshe'; 'Moshe's apartment') (Glinert 1989: 34)

A variation of this pattern is to use morphological affixation to indicate possession. In the examples below, the possessor is shifted from an independent word to an affix on the possessee. In Hebrew ((23a) and (23b)) and the Ethiopian language Tigre (23c), possession by pronouns is indicated by a pronominal affix on the possessee. For instance, in (23a), *sfato* is formed by affixing the morpheme *-o* ('his' / 'its') to the possessee *sfat* ('bank'). This can be contrasted with cases where the possessor is not a prounoun, as in the Hebrew *sfat ha-nahar* (lit. 'bank the-river'; 'the river's bank; Glinert 1989: 29), which uses the same [NOUN1 NOUN2] construction as in (22c); that

[12] It should be noted that there are differences in which of NOUN1 and NOUN2 are the possessor and the possessee. In Yoruba and Hebrew, the possessor is second; in Kobon, the possessor is first.

[13] Actually, this is not the only possibility. Another analysis is that there is no shifting, and the linguistic symbol of the element in question is either a syntactic pattern in conjunction with the overt marker or else just a syntactic pattern (c.f. example (39); also Section 4.1.2.1.3 below).

construction cannot be used if the possessor is a prounoun (c.f. **sfat hu* (lit 'bank it')[14]). In the central North American language Kiowa (23d), an affixation construction can be used regardless of whether the possessor is pronominal or not.

(23) a. sfato
bank.its
'its bank' (Glinert 1989: 29)

b. dodi
uncle.my
'my uncle' (ibid.: 30)

c. səʔli-hom
photograph-3SG
'his photograph' (Croft 2003: 33, citing Raz 1983: 37)

d. nɔ́:- tɔ́: -cègùn
my- brother -dog
'my brother's dog' (ibid., citing Watkins 1984)

There is a similar shifting in the compartmentalization of the information corresponding to subject pronouns in Hebrew that depends on tense. In the case of present tense, a 'full' pronoun is used, such that the pronoun and verb are separate words. In the case of past tense, however, the information corresponding to the pronoun is expressed as an affix on the verb, which is a different structural position. In the Hebrew in (24a) and the English translations for both (24a) and (24b), the information corresponding to 'I' is compartmentalized into a detached pronoun; in (24b), it is shifted (relatively) into a different structural position—that of an affix on the verb.

(24) a. ani choshev she...
I think that
'I think that...'

[14] *-o* is the possessive affix form corresponding to the pronoun *hu* ('he' / 'it').

b. chashavti she...
I.thought that
'I thought that...'

Some constructions that express the equivalent of 'A is B' contain a copula, and others, like those in the Hebrew examples in (25) below, do not and instead use a syntactic arrangement (with no copula). An analogous sort of shifting occurs between these two patterns; the English translations of (25) contain a copula (*am*, *is*, and *am*, respectively), and none of the Hebrew expressions do.

(25) a. achshav ani ha-sho'er
now I the-goalkeeper
'Now I am the goalkeeper' (Glinert 1989: 171)
b. zot ha-siba
that the-reason
'That is the reason' (ibid.)
c. ani be-dirati
I in-apartment.my
'I am in my apartment'

This also depends on tense, like (24) above. In the case of past tense, the copula is used (entailing a relative shift from syntactic arrangement to the overt element[15]), and the pronoun information is encoded as an affix on the copula (entailing a relative shift from 'full' pronoun to verbal affix), as shown in (26).

(26) a. hayiti ha-sho'er
I.was the-goalkeeper
'I was the goalkeeper'
b. hayta ha-siba
that.was the-reason
'That was the reason'

[15] Actually, it is a bit more complicated than this. See the discussion on example (39) in the next subsection.

c. hayiti be-dirati
I.was in-apartment.my
'I was in my apartment'

4.1.2.1.2 Relative separation/compression

Relative separation/compression is variation between the encoding of conceptual material in one element or few elements versus in several elements. The phenomena described by Talmy's *conflation* (discussed above in Section 4.1.1.3) are an important part of this pattern.

I do not intend to suggest that when one expression involves a relative compression of multiple bits of semantic information compared to another, the mental representation of the meaning of that expression necessarily is a composite of those multiple bits of information: there is no reason in principle that the mental representation cannot be atomic (i.e. it can get to a meaning equivalent to one expressed as a composite—c.f. paths to meaning in Chapter 3—but it itself is not necessarily a composite).

One example of this can be seen in a comparison of the Japanese word *genki* and its equivalent expression in English. The word *genki*, a single element[16], corresponds to something like *happy and energetic* in English. The same bit of conceptual material can be compartmentalized into two elements (*happy* and *energetic*)—i.e. relative separation—or into one element (*genki*)—i.e. relative compression.

(27) a. She is <u>happy</u> and <u>energetic</u>
b. kanojo wa genki da
she TOP genki COP

[16] Formally, *genki* is comprised of two morphemes, <*gen*> ('original') and <*ki*> ('spirit'), but is not very semantically analyzable. I am assuming that for the language user it is a single element here.

Similarly, the Lakhota (a central North American language) words in (28) are single elements and correspond to phrases consisting of multiple elements in English (*my / your / his / her* and *mother*).

(28) a. ina
'my mother' (Croft 2003: 33, citing Buechel 1939: 103)
b. nihu
'your mother' (ibid.)
c. huku
'his/her mother' (ibid.)

The Atsuwegi terms like *-iċt* ('into a liquid') and *-cis* ('into a fire') given in (15) above compress (relative to English and many other languages) the information corresponding to 'into' and information about what the object in question is going into a single element.

Relative separation/compression can also be seen between the English and French expressions below in (29). The single element *went* in English (29a) corresponds to the two-word expression *suis allé* in French (29b). This is due to the fact that the French separates tense and aspect, encoding tense in the copula (*suis*) and aspect in the participle (*allé*), while English does not make such a division.

(29) a. I went
b. Je suis allé
I COP.PRS.1SG gone.M.SG

Another sort of relative separation/compression can be seen in polydefinite constructions, like those below, compared to non-polydefinite constructions. Whereas in English, for example, a phrase modified by *the* has only one 'the', in these polydefinite constructions, all of the modifying words have 'the' attached to them. In the Hebrew in (30a), 'program', 'this' and 'new' each receive a 'the'. Similarly, in the Greek in (30b), 'handsome' and 'man' each receive a 'the'. While both Hebrew and Greek have gender and number agreement between the noun and modifying adjective(s), Greek additionally has agreement in the 'the's. This produces a further relative

separation of compartmentalization of the gender and number information from the adjective or noun to both the adjective/noun and the 'the' that attaches to it.

(30) a. ha-tochnit ha-zot ha-chadasha
the-program:F.SG the-this.FSG the-new.FSG
'this new program' (Croft 2003: 36, citing Glinert 1989: 104)

b. o oréos o ántras
the.MSG handsome:MSG the.MSG man:MSG
'the handsome man' [emphatic] (ibid., citing Holton, Mackridge, and Pilippaki-Warburton 1997: 286)

This phenomenon of marking multiple positions versus a single position also occurs in contexts outside the use of determiners. Coordination in Kannada (a southern Indian Dravidian language) marks all of the components being coordinated (31a); by contrast, English, for instance, uses only one 'and'. Whereas the English *between A and B* has only one 'between', the Hebrew construction shown in (31b) attaches a 'between' to both components.

(31) a. Narahariy-u: So:maše:kharan-u: pe: -ge ho:-d-aru
Narahari-and Somashekhara-and market-DAT go-PST-3PL
'Narahari and Somashekhara went to the market' (Haspelmath 2007: 2, citing Sridhar 1990:106)

b. ben A u-ven B
between A and-between B
'between A and B'

A similar multiple-marking pattern is seen in the Chinese *yībiān...yībiān* construction mention in Chapter 3, shown again below in (32). This construction indicates that the two activities marked with *yībiān* are performed in parallel. In the English translations below, the information indicating this parallel action (*while* / *at the same time*) is in a single structural position, by contrast. There is also relative separation/compression between the single-element expression *while* and the multiple-element expression *at the same time*.

(32) a. wǒ xǐhuan yībiān hē chéng zhī yībiān dú bàozhǐ (=(15a) in Chapter 3)
I like *yībiān* drink orange juice *yībiān* read newspaper
'I like to drink orange juice while I read the newspaper.'

b. tā bù něng yībiān xiǎng shìqing yībiān hūxī (= (15b) in Chapter 3)
he not can *yībiān* think thing *yībiān* breathe
'He can't think about something and breathe at the same time.'

In the Chinese construction in (33), the information corresponding to the single English element *into* (again, see the discussion of (39) below) is expressed in two elements, *dào*, which means 'arrive' and functions like 'to', and *li*, which means 'in' (Croft 2003: 76). Thus, the compartmental structure of the Chinese is something like 'I herd sheep arrive in back yard'.

(33) wǒ bǎ yáng gǎn dào hòu yuan li
I OBJ sheep herd to back yard in(side)
'I herded the sheep into the back yard' (Croft 2003: 76, citing Li and Thompson 1981: 400)

Different compartmentalizations of some conceptual material can involve (relatively) pulling apart or fusing different semantic aspects. For example, the distinction between *look* / *listen* and *see* / *hear* is effected in Chinese by 'look' / 'listen' versus 'look-arrive / 'listen-arrive' (34).

(34) a. kàn ('look')
b. tīng ('listen')
c. kàn-dào ('look-arrive'; 'see')
d. tīng-dào ('listen-arrive'; 'hear')

These two semantic aspects can be separated out. For example, 'I can't see' can be expressed as something like 'I look not arrive' (35a), and 'I can't hear' can be expressed as something like 'I listen not arrive' (35b).

(35) a. wǒ kàn bú dào
I look not arrive
'I can't see'
b. wǒ tīng bú dào
I listen not arrive
'I can't hear'

This pulling apart and fusing of semantic aspects can also be seen in the following example, which we also discussed in Chapter 3. In the Chinese in (36b), the thing that the king wants to eat and the thing that the king does eat are separated into different positions/elements (*shénme...shénme*), whereas the wanted thing and the eaten thing are compressed into a single position/element (*whatever*) in English (36a).

(36) a. The king eats <u>whatever</u> he wants. (= (11a) in Chapter 3)
b. guówáng xiǎng chī <u>shénme</u>, jiù chī <u>shénme</u>. (= (11b) in Chapter 3)
king want eat what then eat what
(lit. '(the) king wants to eat something, then (he) eats something')

In relativizing constructions in (37) below, similar separation/compression is seen in the Hebrew and English expressions. In the Hebrew in (37a), the relativized place and the relativization itself are compartmentalized into separate elements (*sham* and *asher*, respectively); in its English translation, however, they are compartmentalized into a single element (*which*). Similarly, in the Hebrew in (b), the relativized person (he who has the wife; *-o*) and the relativization itself (*she*) are compartmentalized into separate elements, whereas the English translation compartmentalizes the corresponding information into a single element (*whose*).

(37) a. ha-adamah <u>asher</u> lukach mi-<u>sham</u> (Genesis 3-23)
the-ground which he.had.been.taken from-there
'the ground from <u>which</u> he had been taken'
b. ha-ish <u>she</u>-raiti et ha-isha shel-<u>o</u>
the-man that-I.saw ACC the-wife of-him
'the man <u>whose</u> wife I saw' (Croft 2003: 229, citing Givón 1979: 183)

Negative concord and polarity are common cases of separation (relative to non-concording negation). Negative concord is essentially the phenomenon of distributing the information pertaining to the negation into several elements. Non-concording negation isolates this information into a single position. The English in (38a) is an example of this. By contrast, the concording negation used in the Spanish in (38b) separates it into two positions. The negative polarity used in (38c) also involves separation of the information, but the negative polarity is not a 'full' negation, and thus this case is somewhat intermediate to (38a) and (38b).

(38) a. I said nothing. None
b. No dije nada. Concord
not I.said nothing
c. I did not say anything. Polarity

Mixtures of separation/compression and shifting are also possible. In Section 4.1.1.3, we discussed Talmy's analysis of the following example, which described the contrast between the English in (39a) and the Spanish in (39b) as a matter of the verb *floated* encoding both movement and manner compared to the Spanish *entró*, which encodes movement only, and *flotando*, which encodes manner only (this is an example of relative separation/compression). But this is not the only possible analysis. In the analysis of Goldberg (1995), in (39a), movement would not be encoded in the verb *floated*, but rather in the construction <VERB into X>; here, *floated*, like the Spanish *flotando,* is held to encode manner only. In this case, there is relative shifting of information about fact and manner of movement between the English and Spanish: the English has movement in the construction and manner in the verb; the Spanish has movement in the main verb and manner in the satellite. This also entails a relative shifting in the structural position of the tense information: tense is attached to the manner-encoding element (*floated*) in English, and to the movement-encoding element (*entró*) in Spanish. There is also relative compression in the English construction, which now additionally contains information about movement. Langacker (2009a) points out that it is also possible for movement to be encoded in both the verb and the construction (i.e. a combination of Talmy's and Goldberg's analyses). So there is relative shifting

depending on whether movement is encoded in the verb or in the construction, and there is relative separation/compression depending on whether it is encoded in a single element (either the verb or the construction), or in multiple elements (both the verb and the construction).

(39) a. The bottle floated into the cave. (Talmy 2000b: 49) (= (11))
b. La botella entró a la cueva flotando. (ibid.)
the bottle moved.in to the cave (floating) (ibid.)

Mixtures of shifting and separation/compression can be seen in Japanese predicate adjective expressions, which in some cases build up the meaning from a compartmental structure that separates tense from the copula and in other cases do not. Japanese adjectives can be broadly divided on the basis of morphology into two groups: *i*-adjectives, which inflect, and *na*-adjectives, which do not. In the case of *i*-adjectives, tense is encoded as an inflection on the adjective, as in (40a); in this case, the copula contains no information related to tense (c.f. **oishii deshita*). This is contrasted with its English equivalent (40b), where tense is encoded entirely in the copula, and not at all in the adjective. This is at once shifting, in that that the tense is transferred between the copula and the adjective, as well as relative separation/compression, in that either the information contained in the copula is split into the copula and adjective or the reverse. The equivalent expression for *na*-adjectives follows the English pattern in this regard, encoding no tense information in the adjective and instead inflecting the same copula used with *i*-adjectives (c.f. (40a)) to encode tense (41a, b), so there is relative separation/compression within Japanese as well, depending on whether the expression uses an *i*- or *na*-adjective.

(40) a. oishi-katta desu
tasty-ADJ.PST COP.NPST
(lit. '(It) is was-tasty')
b. It was tasty

(41) a. meikaku desu
clear COP.NPST
(lit. '(It) is clear')
b. meikaku deshita
clear COP.PST
(lit. '(It) was clear')

A similar sort of mixture of shifting and separation/compression can be seen in a comparison of the Latin and English expressions in (42) below. In the Latin in (42a), pastness is encoded in the participle; in the English in (42b), it is encoded in the copula[17]. So there is shifting of information about pastness between the copula and the participle. Additionally, there is relative separation/compression in two places. The English *was* contains both the copular meaning and pastness; these aspects are separated out in the Latin. Also, the single Latin word *sum*, which contains information about person and number of its subject and thereby incorporates the information encoded by *I*, corresponds to two words (the pronoun and copula, *I was*) in English; so there is separation/compression between *I was* and *sum*[18].

(42) a. damnatus sum
condemned.PST COP.1SG.PRS
(lit. 'I am was-condemned')
b. I was condemned

Another kind of mixture of separation/compression and shifting can be seen in the following examples. Verbs in the Oklahoman language Caddo include a morpheme that indicates the sort of thing that the patient is, as in (43) below (Talmy 2000b: 112-113). Between the English and the Caddo, there is relative separation/compression for 'milk' /

[17] In contrast to the case of the Japanese *desu* in (40), it is not accurate to say that *sum* here does not carry tense. Unlike *desu* there, *sum* can inflect for tense, as in *damatus fui* (lit. 'I was was-condemned'; 'I had been condemned').

[18] It should be noted that the pronoun can be included—e.g. *ego damnatus sum* is also possible—or omitted here.

‘salt’ versus ‘milk’ + ‘liquid’ / ‘salt’ + ‘powder’ and shifting of structural position of the ‘liquid’ / ‘powder’ bit between the noun (English) and the verb (Caddo).

(43) a. cú·cuʔ *kan*-yi-daʔk-ah
milk liquid-find-PST
lit. ‘He liquid-found the milk’; ‘He found the milk’ (Talmy 2000b: 112)
b. widiš *dáʔn*-yi-daʔk-ah
salt powder-find-PST
lit. ‘He powder-found the salt’; ‘He found the salt’ (ibid.)
c. *dáʔn*-yi-daʔk-ah
powder-find-PST
lit. ‘He powder-found it’; ‘He found it (something powdery)’ (ibid.: 113)

Another case can be seen in (44). The Canadian language Cree (44a) marks the possessee with a possessive affix that carries pronominal information about the possessor (in a form like ‘John his-canoe’). Compared to the English translation, there is shifting of possession information from ‘John’ (English) to ‘canoe’ (Cree); also, there is separation of information relating to the possessor into *cān* (‘John’) and *o-* (‘his’). One way of expressing possession in Hebrew (44b) uses a similar compartmental structure to the Cree in (44a), except there is a further relative separation of possession information into the possessive affix and *shel* (‘of’).

(44) a. cān o-cīmān
John 3SG.POSS-canoe
‘John’s canoe’ (Dryer 2007: 178, citing Ellis 1983)
b. dirato shel moshe
apartment.his of Moshe
‘Moshe’s apartment’ (Glinert 1989: 34)

4.1.2.1.3 Indexation vs. word position

In Section 2.2.1.2 (What counts as an element?), I defined a construction as ‘a syntactic pattern that carries a particular meaning’, and I pointed out in footnote 32

there that (depending on the language) syntactic patterns can be established by word position and by patterns of indexation: constructions can be instantiated through patterns of word position and patterns of indexation. Indexation is the use of marking (case, etc.) to indicate the syntactic role of words, distinct from word position. I elaborate on this here. This is an important part of how language functions, and it is related to compartmentalization because it has to do with the structuring of information in expressions.

Consider, for example, the linking of an adjective and the noun it modifies in English. This is done purely through word position: there is a fixed word position pattern (*angry cat* vs. **cat angry*), and absolutely no indexation is employed (nor are other means like indicating the attachment with a particle).

Spanish uses both word position and indexation to establish the same linking. Like English, the word position is fixed (the reversed order is ungrammatical). But there is also agreement in number and gender in the noun and the adjective. The number and gender marking on the noun and adjective functions as an indexation of those elements that, together with word position, establishes the linking.

(45) a. gato enojado (c.f. *enojado gato)
cat.M.SG angry.M.SG
'angry (male) cat'

b. gata enojada (c.f. *enojada gata)
cat.F.SG angry.F.SG
'angry (female) cat'

c. gatos enojados (c.f. *enojados gatos)
cat.M.PL angry.M.PL
'angry cats'[19]

[19] In Spanish, masculine is the "default" gender that is used for mixed-gender groups; *gatos enojados* can therefore mean 'angry cats' (with no specification of gender), 'angry (male and female) cats', or 'angry male cats'.

d. gatas enojadas (c.f. *enojadas gatas)
cat.F.PL angry.F.PL
'angry female cats'

Latin, by contrast, uses pure indexation to establish the linking. Unlike English and Spanish, the word order does not matter for establishing the linking, as seen in (46a) and (46b). In fact, the noun and adjective do not even have to be adjacent at all, as shown in (46c), where they are separated. Word position in Latin is extremely free; all four of the words in (46c) can go in any of the four positions. Word position itself has absolutely nothing to do with establishing any of the linkings in (46c) ('angry' and 'cat'; 'cat' and bit'; 'bit' and 'boy'; perhaps also 'cat' and 'boy'). Instead, indexation is used. Indexation in Latin nouns and adjectives is achieved by marking case, number, and gender. Agreement in this marking is what establishes that *īrātus* and *fēlis* are linked (and that, for example, it is not the boy who is angry). Thus, the syntactic pattern that is the construction here resides in this indexation pattern (and not in word position).

(46) a. fēlis īrātus
cat.NOM.M.SG angry.NOM.M.SG
'angry cat'

b. īrātus fēlis
angry.NOM.M.SG cat.NOM.M.SG
'angry cat'

c. īrātus momordit puerum fēlis.
angry.NOM.M.SG bit.3.SG boy.ACC.M.SG cat.NOM.M.SG
'(The) angry cat bit (the) boy.'

The other linkings are likewise established by indexation: that 'cat' is the subject of 'bit' is established by its nominative case marking[20] and number and person

[20] This is actually not always the case in Latin. The subject of absolute clauses takes the ablative case, where it is linked to the verb through indexation, as in:

rē dēcretā decessit.
matter:F.ABL.SG decided.F.ABL.SG leave.PRF.3SG
'The matter having been decided, (he / she) left.'

agreement with ‘bit’, and that ‘boy’ is the object of ‘bit’ is established by its accusative case marking. Any changes in agreement or in case would destroy the linking.

Another example of case marking patterns establishing the construction can be seen in the following example from the Australian language Thalanyji (Evans and Levinson 2009: 441, citing Austin 1995: 372). Here, the dog belongs to the woman, but (unlike in English) there is no fixed linear position of the words that corresponds to the possession relationship. The linking between ‘dog’ and ‘woman’ is established instead by the case marking.

(47) a. Kupuju-lu kaparla-nha yanga-lkin wartirra-ku-nha
child-ERG dog-ACC chase-PRS woman-DAT-ACC
‘The child chases the woman’s dog.’

Now, some constructions are rather tricky, since they are not fully specified by indexation patterns or by patterns in word position. One example is the attachment of genitives in Latin. The attachment of *ultima* (‘last’) to *aetas* (‘age’) is achieved by case agreement, just as in (46) above. Latin marks only the possessor, and not the possessee, so the sort of indexation described above (which requires marking on all words that are to be linked) does not apply[21]. There are two genitives here: *Cumaei* (‘of Cumae’) and *carminis* (‘of (the) song’); they are indistinguished by case marking and are in totally different linear positions with respect to the words they modify (i.e. *carminis* and *aetas*, respectively). There is nothing in either the word position or the indexation that indicate what they are attached to. But they are nevertheless attached. (Only case marking is shown here, because that is what is relevant to my point. There is no other marking that indicates the linking.)

[21] This system differs from that of languages like Turkish, which do mark both the possessor and the possessee:

Ahmed-in o -u
Ahmet-GEN son-3SG.POSSESSED
‘Ahmet’s son’ (Dreyer 2007: 185)

(48) ultima Cumaei venit iam carminis aetas
last.NOM Cumae.GEN came now song.GEN age.NOM
'The last age of the Cumaean song has now arrived' (Virgil; cited in Evans and Levinson 2009: 441)

In these cases, the construction still exists, of course (as a matter of necessity), but it is very much underspecified by the linguistic form. The only way the expression can be interpreted (with regard to what belongs to what) to have the meaning that it does is on the basis of the fact that interpretation is what makes sense. But for the problem of expressing thought in language, these cases are not so different from the 'normal' ones: as far as the language user is concerned, there is a linguistic element (here, one that indicates a kind of possession relationship between two other elements, whose linguistic symbol is genitive case marking on the element that does the possessing) that can be used to express some meaning.

4.1.2.1.4 Fundamental Reformulation

The comparison of the examples discussed above involved roughly similar information between the expressions. However, in other cases, a particular effective meaning can be gotten at with expressions that do not share information. This is what I refer to as *fundamental reformulation*. In any comparison of multiple expressions, the information they contain will be different to some extent; in some cases, the information will be more different, and in other cases less. Overall, similarity in the information is a matter of degree; the examples included in this section contain strikingly different information and organization of that information.

The example first brought up in Chapter 3 of 'necessary' being obtained by 'absolute importance' (*hitsuyou*; Japanese) or 'not yielded' (*necesssus*; Latin)

demonstrates the obtaining of similar meaning with unrelated information quite clearly.[22]

The same basic principle can also be seen in (49) below. The English *I'm thirsty* (49a) and the Spanish *Tengo sed* (49b) contain similar information, although they structure it differently. By contrast, the Japanese equivalent *nodo ga kawaita* (49c) does not contain similar information: the conceptual material depicted is of dryness rather than of thirst[23].

(49) a. I'm thirsty

b. Tengo sed
I.have thirst
(lit. 'I-have thirst')

c. nodo ga kawai-ta
throat SBJ become.dry-PFV
(lit. 'throat is dry')

Fundamental reformulation can also be seen between the English and Japanese expressions below. Japanese expresses 'some people...' using a compartmental structure like 'there are also people who...'. These expressions use fundamentally different information to produce similar meanings.[24]

[22] Both of these exhibit relative separation compared to *necessary*, but that does not mean that their mental representations share content (either with each other's or with that of *necessary*). Pulling apart semantic aspects from an atomic mental representation into a composite one results in *different* mentally-represented content. This is part of why I note that the information encoded in different expressions is always different to some degree and why I mentioned in the introduction to Section 4.1.2.1 that these patterns (shifting, separation/compression, and fundamental reformulation) are prototypical in essence and not wholly discrete.

[23] Also, in the Japanese, 'dry' is expressed in perfective aspect; this is totally absent in the English and Spanish expressions.

[24] Also compare the involvement of the additive in the generation of 'some people' here with the generation of 'anyone' from 'who' + additive seen in (18).

(50) a. Some people think the earth is flat.
b. chikyuu wa taira to omotteiru hito mo iru
earth TOP flat COMP think person ADDITIVE be.NPST
(lit. 'There are also people who think the earth is flat.')

Japanese uses the same structure for statements (51a) and questions (51b) of ordinal numbers. However, English, in fact, does not have a construction that permits asking ordinal numbers[25], and instead requires a fundamental reformulation, such as that in (51c).

(51) a. obama wa 44-banme no daitouryou desu.
Obama TOP 44-ORDINAL GEN president COP.NPST
(lit. 'Obama is the 44th president.')
b. toranpu wa nan-banme no daitouryou desu ka?
Trump TOP what-ORDINAL GEN president COP.NPST Q
(lit. 'Trump is the what-th president?') (Akira Machida, p.c.)
c. What number president is Trump?

Similarly, English structures meaning analogously in expressions that use *more* (52a) and *less* (53a); in the case of 'more', Japanese uses a compartmentalization quite similar to that of English (52b), but in the case of 'less', Japanese requires a fundamental reformulation, along the lines of (53b) (replacing 'less' with 'so much' and a negative imperative[26]). This is because, while Japanese *motto* is analogous to *more*, Japanese lacks a lexical analog of *less*, and so the meaning must be structured differently.

[25] The literal translation of (51b), *Trump is the what-th president?*, which 'forces' the same structure used for stating ordinal numbers, is probably understandable, however.

[26] In comparison to the literal translation of (53b), there is also relative separation of 'much' into the degree (*sonnani*) and the amount (*ooku*) (a 'strict' literal translation of the version without *ooku* would be 'Do not put-in to such a degree.'); this relative separation/compression does not exist between (52a) and (52b).

(52) a. Put in more next time.
b. jikai, motto irete.
next.time more put.in.IMP
(lit. ‘Put in more next time.’)

(53) a. Put in less next time.
b. jikai, sonnani (ooku) irenaide.
next.time so.much (much) put.in.NEG.IMP
(lit. ‘Do not put in so much next time.’)

The English expression *A is special in the same way B is* (54a) can be roughly paraphrased as ‘A is special in a particular way. B also has that specialness.’; the Japanese in (54b) approximates this fairly closely. The same basic compartmental structure of (54a) can be used in *A is special in a way that B is not* (55b), which is roughly paraphrasable as ‘A is special in a particular way. B lacks that specialness.’; in this case, however, Japanese cannot use an analogous structure. This is because in order to use that structure to express a meaning like that of (55a), a negation of *you ni* (to effect a meaning like ‘unlike’, as opposed to the ‘like’ in (54b)) would be necessary (55b); however this is not possible in Japanese. As a result, a fundamental reformulation, such as that in (55c), is required. This involves a relative shifting of the negation from ‘like’ to ‘is in’, reconceptualizing ‘special’ as ‘specialness’, and reconceptualizing the link between ‘A’ ‘special(ness)’ as one of possession, rather than copular linkage. Other aspects of differences in meaning attendant to this reformulation will be discussed in Section 4.1.2.2 (Interdependencies) below.

(54) a. A is special in the (same) way B is.
b. A wa B no <u>you</u> ni tokubetsu da.
A TOP B GEN like ADV special COP
(lit. ‘A is special <u>like</u> B’)

(55) a. A is special in a way that B is not.
b. A wa B no [¬ <u>you</u>] ni tokubetsu da.
A TOP B GEN NEG.like ADV special COP
(lit. ‘A is special <u>unlike</u> B’)

c. A no tokubetsusa wa B ni nai.
A GEN special.NMLZ TOP B in be.not
(lit. 'A's specialness is not in B.')

Comrie and Thompson (2007: 358-359) point out that 'the enemy's destruction of the city' cannot be expressed in Russian using a compartmental structure similar to the English (56a). Using possession to link the nouns, as in (56b) and (56c), results in different meanings; it cannot indicate that the enemy is the agent of the destruction. The nearest equivalent of (56a) in Russian would involve reformulating the conceptual material, such that the relationship between 'the enemy' and 'destruction' is expressed by something like 'by' (56d); this depicts a different conceptual structure from the English in (56a).

(56) a. the enemy's destruction of the city (ibid.: 358)
b. razrušenie goroda vraga
destruction of.city of.enemy
'the destruction of the enemy's city' (ibid.)
c. razrušenie vraga goroda
destruction of.enemy of.city
'the destruction of the city's enemy' (ibid.)
d. razrušenie goroda vragom
destruction of.city by.enemy
'the destruction of the city by the enemy' (ibid.: 359)

Above, we discussed relative separation/compression between *whatever* and *shénme...shénme* in (36), reproduced below as (57). Another difference, not discussed there, is that the Chinese (57b) uses a sort of conditional, whereas English does not. As described above, the compartmental structure in the Chinese separates what English depicts as a single event into two events (a wanting event and an eating event). The linking of these two events is achieved by a (quasi-)conditional construction; this sort of conceptualization is not depicted by (57a).

(57) a. The king eats whatever he wants. (= (36a), (11a) in Chapter 3)
b. guówáng xiǎng chī shénme, jiù chī shénme. (= (36a), (11b) in Chapter 3)
king want eat what then eat what
(lit. '(the) king wants to eat something, then (he) eats something')

4.1.2.2 Interdependencies in compartmental structure

Another important point concerning compartmental structure and compartmentalization is that different parts of an expression interact with each other in constructing meaning. Compartmentalizing part of the thought in a particular fashion places restrictions on how other parts of the thought can be compartmentalized in order to coherently produce a suitable meaning. The interrelationships between different potential parts of an expression that cause these restrictions are what I refer to as *interdependencies*. As a consequence of these interdependencies, changing one part of an expression often necessitates changing another. This means that the compartmentalization of the thought must be made in consideration of the interactions between its various parts. Some of this relates to whether a particular bit of conceptual material is covered by a potential part of the expression or not (e.g. a language user who would be satisfied by *genki* as an expression of their thought may feel that *happy* alone is insufficient, because it does not cover the 'energetic' aspect)[27], some of it depends on (combinatorial) meaning construction (e.g. *not yielded* and *absolute importance* can construct a meaning like 'necessary', but *not important* and *absolutely yielded* do not), some of it concerns incompatibilities between different potential local compartmentalizations, and some of it is set when other parts of the surrounding linguistic context are fixed. The discussion below will focus on these last two aspects (which are not really separate from each other), as the first is relatively straightforward and the mechanics behind the second were already discussed in Chapter 3.

Yamanashi (2000: 246-253) points out that the Japanese expressions (58a) and (58b) can function as paraphrases of each other; however, while *zenzen* ('at all') (a

[27] This includes cases where a potential part of the expression would say too much (i.e. include unwanted semantic aspects).

negative polarity item) can be added to (58a) for emphasis (58c), it cannot be added to (58b): the indirect quasi-negation of 'outside' cannot interact properly with the negative polarity of *zenzen* (58d). The English translations exhibit the same phenomenon. If part of the thought is compartmentalized as *zenzen* / *at all*, the part of the thought relating to the 'negation' (in a broad sense) has to be compartmentalized as in (58a) and not (58b).

(58) a. kore wa watashi no senmon dewanai.
this TOP 1 GEN specialty COP.NEG
(lit. 'This is not my specialty.')

b. kore wa watashi no senmon-gai dearu.
this TOP 1 GEN specialty-outside COP
(lit. 'This is outside my specialty.')

c. kore wa zenzen watashi no senmon dewanai.
this TOP at.all 1 GEN specialty COP.NEG
(lit. 'This is not my specialty at all.')

d.* kore wa zenzen watashi no senmon-gai dearu.
this TOP at.all 1 GEN specialty-outside COP
(lit. *'This is outside my specialty at all.')

Similarly, the construction in (59a) can be passivized, as in (59c), and it can be modified with *try*, as in (59b), but it cannot be both passivized and modified with *try* (59d), as these two constructions are incompatible with each other. A language user can compartmentalize part of their thought as *trying to* [convey], and they can compartmentalize another part of their thought as *being conveyed*, but they cannot express both of these compartmentalizations.

(59) a. The author is conveying their perspective to the reader.
b. The author is trying to convey their perspective to the reader.
c. The perspective of the author is being conveyed to the reader.
d.* The perspective of the author is being tried to be conveyed to the reader.

This is sometimes done anyway, however, since often there is no other acceptable way to express the thought. Turner (1998) discusses the attested utterance *I was hit by*

the judge; I was tried to be hit by the umpire, which was said by a coach describing an incident at the 1988 Olympics; Turner remarks that if the expression were edited to be well-formed (e.g. as in *I was assaulted by the umpire*), it would probably be deficient in expressing the desired meaning in some aspect.[28]

The expressions in (60) are all quite similar in meaning. If a language user determines that one of *possible* or *can* is a better fit for part of their thought, they must change other parts of the expression as well. A simple exchange does not work (e.g. **It is can (to) imagine that...*). One way to restructure (60a) for the use of *can* is to add a virtual imaginer (*one*), as in (60b); another is to make *imagine* passive, as in (60c).

(60) a. It is possible to imagine that...
b. One can imagine that...
c. It can be imagined that...

As mentioned in Chapter 3, Tübatulabal has two markers that distinguish between continued and interrupted simultaneous action. Those examples are reproduced below as (61a) and (61b). In the distinction between whether one action continues through the occurrence of the other or is interrupted by it, these markers are similar to *while* and *until* in English. Now, both of the Tübatulabal markers are used on the verb of the subordinate clause (i.e. it is the action that is either continued or interrupted that is subordinate); however, switching between *while* and *until* requires alternating which material goes in the main clause and which material goes in the subordinate clause. With *while* (corresponding to continued simultaneity), 'the woman was eating' goes in the subordinate clause (like Tübatulabal) (61c); in this case, having it in the main clause is not possible. Conversely, it is not possible to express the interruption (interrupted

[28] In cases like these, whether a particular combination of compartmentalizations is expressible or not (i.e. whether the resulting expression is acceptable or not) can vary between individuals, depending on their analysis of the relevant linguistic structures. For example, Pylkkänen and McElree (2006) give *The book was begun to be written by the author* as an example of a deviant expression, but De Almeida and Dwividi comment (2008: 319 note 13) that they consider it well-formed (c.f. *The book has begun to be written*, which is probably acceptable to most speakers).

simultaneity) in English (this has to use *until*) and have 'the woman was eating' in the subordinate clause, unlike Tübatulabal (61d).

(61) a. kó:imí ánaŋ-áŋ iŋgím tá:twál (= (10b) in Chapter 3)
woman cry-SIM1 came man
'While the woman was crying[SIM1], the man came.' (Corbett 2007: 309)

b. kó:imí tïka-káŋ apá'agín tá:twál (= (10c) in Chapter 3)
woman eat-SIM2 hit man
'The man hit the woman when the woman was eating[SIM2] [and as a result her eating was interrupted].' (ibid.)

c. The man came in while the woman was eating

d. The woman was eating until the man hit her

The Hebrew examples discussed in (25) and (26) above, reproduced below as (62) and (63), exhibit interdependencies between tense and whether the copular joining is effected by a syntactic arrangement (62)—this is used in the case of present tense—or an overt copula (63)—this is used in past tense. Tense also affects the compartmentalization of the subject: in the case of present tense, it is expressed as a pronoun (62); in the case of past tense, it is expressed as an affix on the verb (63).

(62) a. achshav ani ha-sho'er (= (25a))
now I the-goalkeeper
'Now I am the goalkeeper' (Glinert 1989: 171)

b. zot ha-siba (= (25b))
that the-reason
'That is the reason' (ibid.)

c. ani be-dirati (= (25c))
I in-apartment.my
'I am in my apartment'

(63) a. hayiti ha-sho'er (= (26a))
I.was the-goalkeeper
'I was the goalkeeper'

b. hayta ha-siba (= (26b))
that.was the-reason
'That was the reason'

c. hayiti be-dirati (= (26c))
I.was in-apartment.my
'I was in my apartment'

Another kind of interdependency exists in Rif Berber expressions where an adjective modifies a noun; this interdependency is between definiteness of the noun and whether the modification is done directly or in a relative clause: if the noun phrase is definite, the modification is done directly (64a), but if it is indefinite, the adjective must go in a relative clause (64b) (Dryer 2007: 174). Thus, the Rif Berber equivalent of *the small boys* is something like 'the small boys', but that of *large pieces of wood* is something like 'wood that is large'.

(64) a. ramən *imeẓẓyanən*
boy.PL small
'the small boys' (Dryer 2007: 174, citing Kossman 2000)

b. iqəššuḍən [d iməqqw anən]
wood COP large
'large pieces of wood' (ibid.)
(lit. 'wood that is large')

Above, I mentioned that in Japanese, the compartmentalization of tense in Japanese predicate adjective expressions depends on the type of adjective (examples (40) and (41)). If an *i*-adjective is used, tense is carried by the adjective (and not the copula); if a *na*-adjective is used, tense is carried by the copula (and not the adjective). Some *i*- and *na*-adjectives can be quite close in meaning to each other; an example is given below in (65): both *subarashii* (an *i*-adjective) and *suteki* (a *na*-adjective) can be translated as 'wonderful'. Choosing one or the other also affects how the tense information must be compartmentalized.

(65) a. It was wonderful.
b. subarashikatta desu.
wonderful.PST COP.NPST
(lit. '(It) is was-wonderful.')
c. suteki deshita.
wonderful COP.PST
(lit. '(It) was wonderful.')

I also pointed out earlier that expressing 'A is special in a way that B is not' (66a) in Japanese requires a fundamental reformulation compared to the English in (66a) and to the case where both A and B are special (66b); one way to express it is as 'A's specialness is not in B' (66c). Now, conducting this reformulation affects several aspects of the meaning of the expression. One notable aspect that is affected is that while *A is special in a way that B is not* (66a) directly asserts that A is special, 'A's specialness is not in B' (66c) only permits implying it. Another concerns the expression of fact of specialness versus quality of specialness: *A is special in a way that B is not* expresses simultaneously both quality and fact (i.e. both the fact that it is special and the particular way in which it is special are not shared by B) 'A's specialness is not in B' is ambiguous between fact and quality (i.e. either A's particular kind of specialness is not in B or B lacks specialness altogether).

(66) a. A is special in a way that B is not. (= (55a))
b. A wa B no [¬ you] ni tokubetsu da. (= (55b))
A TOP B gen NEG.like ADV special COP
(lit. 'A is special unlike B')
c. A no tokubetsusa wa B ni nai. (= (55c))
A GEN special.NMLZ TOP B in be.not
(lit. 'A's specialness is not in B.')

So when a language user compartmentalizes a thought, multiple bits of conceptual content are often relevant at once; how the meanings of the different elements of the expression interact and work together must be taken into account[29].

4.2 The selection of linguistic elements

The other facet of constructing an arrangement of elements to express a thought is selecting the elements that are to go into it: in order to produce an expression, one must obviously choose the appropriate words etc. that make it up. Categorization has been identified as a means by which this is done.

In Section 4.2.1, I explicate the notion of categorization, first from the perspective of psychology, and then from the perspective of cognitive linguistics, focusing on its complexity. In Section 4.2.2, I discuss results of psycholinguistic research that have determined that access of elements [30] can be affected by priming and involves associative mechanisms. In Section 4.2.3, I explain that categorization is insufficient as an account of the selection of elements and that the more comprehensive notion of compartmentalization answers its lack.

4.2.1 Categorization

It is widely accepted that the selection of elements involves categorization. Simplistically, categorization can be thought of taking one concept and finding a closest

[29] I do not mean to imply that compartmentalization must be done of the entire thought at once in a single step. See Section 5.3 for a detailed discussion of how compartmentalization is incorporated into the process of expressing thought in language; this is elaborated on there.

[30] More commonly-used terms are *lexical access* and *lexical selection*; I refer instead to *access of elements* and *selection of elements* here because not all linguistic elements are lexical—in fact, some of the phenomena discussed below relate to non-lexical linguistic structures (i.e. constructions).

matching concept from some other store of concepts, in our case from the set of all concepts in the target language[31]. But it is not always so simple.

Hampton (2017) gives an overview of psychological research on categorization and describes some of the complexities involved. Essentially, when comparing some item with some category, the central dimension of how well it matches the category (i.e. family resemblance) is not the only factor in the categorizations that people make. Items that occur with higher frequency or are more familiar are more surely categorized; rarer items tend to be judged to be less typical of a category (also c.f. Barsalou 2017). Ideals are an additional complicating factor. 'Good' items ('good' in two senses—good at exemplifying the category and good in terms of evaluative judgment) tend to be considered more typical of a category, even though they tend to be more extreme in their properties and therefore deviate from the 'center' of the category. An ideal winter coat is not the same as an average winter coat. Notably, similarity alone does not determine the appropriateness of a categorization: as Hampton explains (2017: 108-109), a bat is not an atypical bird—it is not a bird at all, despite having some properties that are more prototypical of birds than some actual birds (e.g. bats fly, but ostriches do not). Typicality is thus distinguishable from category membership. Other effects show up in categorizations related to combinations of concepts. Typicality with respect to the component categories does not translate straightforwardly to typicality in the combined category. For example, a guppy is more prototypical of the category [pet fish] than of either that of [pet] or [fish] (Osherson and Smith 1981; referenced also in Ran and Duimering 2010: 67; Hampton 2017: 109). Similarly, Jönsson (2015) found that when participants viewed a video of a person smoking in a typical fashion and simultaneously walking in an atypical fashion, the person was without exception judged to be [smoking] and judged to be [walking] by only a minority of participants, but judged to be [smoking and walking] by 70% of participants.

Other aspects of the complexity of categorization are elucidated by Langacker (2009a). He describes a typology of categorization based on different relationships between the *target* concept to be categorized and a pre-established category (called the

[31] That the language user knows, as they know them. The 'target language' here is what is inside the mind of the language user (c.f. Langacker 2008: 217).

standard) against which it is to be categorized. He identifies three principle types: *full recognition*, *partial recognition*, and *complex recognition*. He diagrams them in Figure 8 below (the target is abbreviated as T and the standard as S). Full recognition (a) is the case where the target concept is directly apprehended as the standard. In cases of full recognition, the target and standard are "co-extensive" (i.e. the boundaries of their domains are identical—they overlap completely) (this is what is meant by the arrows), but the target can be more detailed or elaborate than the target. In the case of partial recognition (b), the target and standard only partially overlap. Langacker gives three subtypes of partial recognition. When the partial recognition is *contrastive* (i), the target is almost the same as the standard, but has qualitative differences in some features. When it is *subtractive* (ii), the standard has features that are absent in the target.[32] When it is *augmentative* (iii), the standard contains features absent in the target, and as a result of apprehending the target as the standard, some features originally not present in the target are added to the target; this produces a new target (T′). Complex recognition (c) is the case where the target is perceived as a combination of multiple standards.

[32] It is not entirely clear to me what the difference between the contrastive and subtractive cases is intended to be. The diagram suggests that in the contrastive case, some of the content of the target is missing in the standard, and some of the content of the standard is not contained in the target, whereas in the subtractive case, the target corresponds to a subsection of the standard. However, for the contrastive case, Langacker gives the example of someone's mother dyeing their hair and still being recognized as the same mother; for the subtractive case, he gives the example of a cat with a missing tail. These examples imply that the distinction is instead between whether a feature exists or not (e.g. the tail present or missing) versus what the value for a feature is (e.g. the color of one's mother's hair).

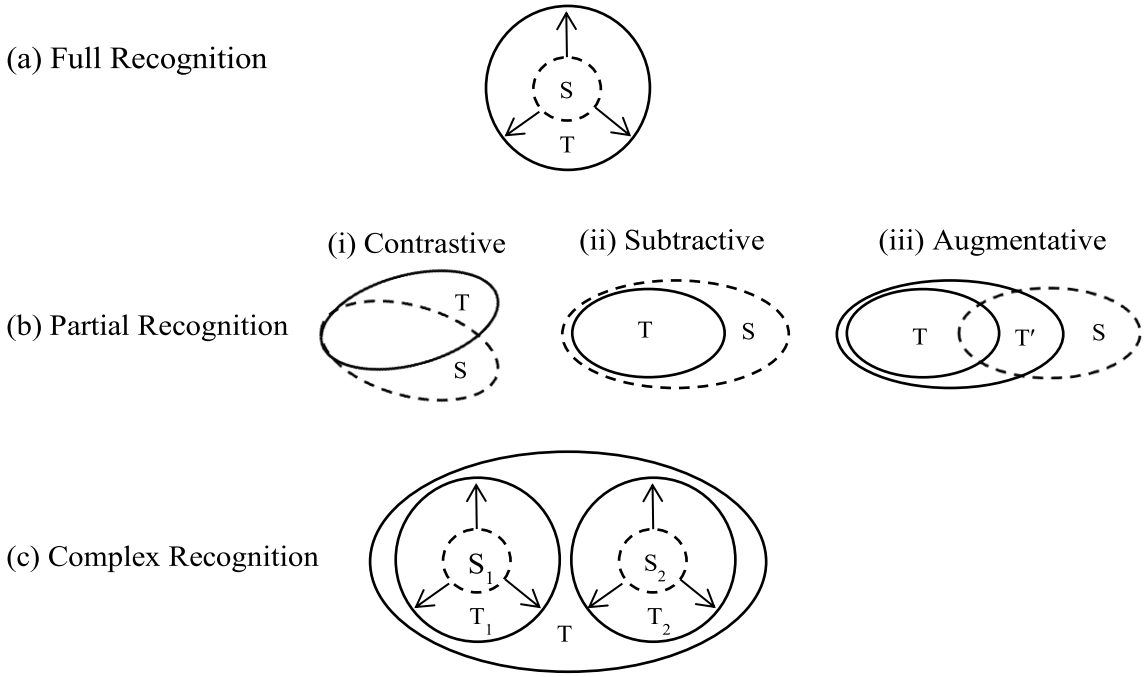

Figure 8: Langacker's typology of categorization (2009a: 228)

So there are many different ways that an item to be categorized can match and diverge from a particular category.

There are many ways in which the fit between language and thought is imperfect (some were discussed earlier in Section 2.2.1.3 (Granularity of linguistic meaning)). An expression can say too much or too little, or be too specific or too vague; it can deviate in simply being conceptually different or in taking a distinct path to meaning not represented in the original language-independent conceptualization. The problem that must be solved in language production is to find an expression that can be considered to not deviate too much.

4.2.2 Accessing elements

Another aspect of the issue of the selection of linguistic elements is the cognition of accessing the elements. Psycholinguistic research has discovered two principle related aspects of the mechanics of this cognition. One is that the access of elements is

subject to priming effects: if a language user is exposed to a particular linguistic structure, they process that structure and related ones faster and are more likely to use that structure (or a related one) in language production. The other is that the access of elements is associative: accessing one element automatically activates others that are related to it. This automatic access is often referred to as spreading of activation. A large part of the empirical basis for the understanding that the access of elements is associative is the data on priming effects; the cognitive mechanics underlying associative access and priming effects are probably basically the same, although they are distinct phenomenologically (e.g. associative access can occur without external priming).

Important evidence for associative access was discovered by Meyer and Schvaneveldt (1971). They used an experimental technique of having participants determine whether a presented string of letters is a word or not (known as a *lexical decision* task); here, they presented two (potential) words simultaneously and had the participants determine, as quickly and accurately as possible, whether they were words or not[33]. Some of the pairs of words they used were semantically associated with each other (e.g. *bread* and *butter* or *nurse* and *doctor*). They found that reaction times for the lexical decision task were shorter when the words in the stimulus pair were associated (e.g. *bread–butter* or *nurse–doctor*) compared to when they were not associated (e.g. *bread–doctor* or *nurse–butter*).

Priming effects in lexical access have been very extensively researched. In addition to facilitation effects like those found by Meyer and Schvaneveldt, interference effects also exist; here, certain kinds of associations actually retard lexical access. In certain circumstances, distractor words with related meanings facilitate access of the target word, and in other cases inhibit it. If a word denotes a competing category (e.g.

[33] The pairs of strings Meyer and Schvaneveldt presented came in one of three combinations: a pair of words, a pair of non-words, or one word and one non-word. They gave the participants two tasks. In one task, they had participants judge whether both strings were words or not (i.e. a pair consisting of non-words and a pair consisting of a word and a non-word would both be 'no'). In the other task, they had participants judge whether the status as words or non-words was the same for both strings (i.e. a pair of words or of non-words would be 'same'; a word and a non-word would be 'different').

'dog' vs. 'cat'), picture naming can be slowed—unless the task is to name the superordinate category (e.g. 'animal'), in which case it is faster instead; presentation of *frisbee* delays the naming of a picture of a ball, but *game* quickens it (Roelofs 1992; Meyer 1996; Cutting and Ferreira 1999; Costa et al. 2005; Dell et al. 2014; etc.). Representations of phonological forms have also been found to influence word retrieval: for example, in picture-naming tasks, when a distractor word phonologically similar to the target word is presented together with the stimulus, response times are affected (e.g. Meyer 1996; Cutting and Ferreira 1999; Damian and Martin 1999; see also Goldrick 2014).[34]

Bock and Loebell (1990) found that priming effects also exist for constructions. They presented a prime that was of a structure like (in one experiment) either *The wealthy widow gave an old Mercedes to the church* or *The wealthy widow sold the church an old Mercedes* or (in another experiment) either *The construction worker was hit by the bulldozer* or *The construction worker drove the bulldozer*. They had participants repeat the prime sentence aloud and then describe a scene shown in a picture. They found a tendency for the construction used by the participants to describe the scene to align with the structure of the prime (e.g. participants more frequently used passives when the prime was passive than when it was active). This effect is called *structural* or *syntactic priming*. It has been found to occur across languages. Loebell and Bock (2003) tested German-English bilinguals using similar primes in one language (e.g. *The lawyer sent his client the contract* vs. *The lawyer sent the contract to his client* or the German analog) and had participants describe the scene depicted in a picture in

[34] The effect for both semantically- and phonologically- related words also depends on variation of the timing (on the order of a few hundred milliseconds) of the presentation of the distractor word.

the other language and found the same effect[35]. Hartsuiker et al. (2004) tested Spanish-English bilinguals using active or passive primes in Spanish (e.g. *El taxi persigue el camion* ('The taxi chases the truck') vs. *El camion es perseguido por el taxi* ('The truck is chased by the taxi') and found that participants tended to use the corresponding construction in English when describing a picture. Moreover, language users can prime themselves—i.e. using a particular construction makes it more accessible and relatively more likely that it will be used again (Jacobs et al. 2019).

Additionally, linguistic structures that occur with higher frequency in general use tend to be more easily accessed, reflecting higher mental accessibility of those structures. Oldfield and Wingfield (1965) found that when participants had to name objects presented to them, objects whose name is more frequently occurring linguistically were named faster. Alario et al. (2002) investigated this phenomenon for adjective + noun expressions. They had participants name pictures (e.g. a blue kite); the response time of the participants correlated with the frequency of both the adjective and the noun. Arnon et al. (2010) had participants judge whether a four-word phrase presented to them was acceptable or not (e.g. *I saw the man* vs. *I saw man the*); responses were faster for frequent phrases (e.g. *don't have to worry*)[36]. Janssen and Barber (2012) had participants produce noun + adjective, noun + noun, or determiner + noun + adjective expressions in Spanish and French based on stimulus pictures, and found that the speed of response depended on the frequency of the phrase as a whole (as distinct from that of the object name).

---

[35] Loebell and Bock also tested cross-linguistic priming of passive constructions and found no effect. However, they note that German and English passives are not as structurally analogous as dative and double-object constructions (as in the lawyer example): passives in German can be formed with the main verb at the end of the sentence (e.g. *Die Böden werden täglich von dem Hausmeister gereinigt*, glossed as 'The floors are daily by the janitor cleaned'), whereas with dative and double-object constructions, word-for-word glosses are identical between the languages (e.g. *The boy sent his pen-pal a letter* / *Der Junge schickte seinem Brieffreund einen Brief* vs. *The boy sent a letter to his pen-pal* / *Der Junge schickte einen Brief an seinen Brieffreund*) (pp. 795-797). They tested within-language priming of German passives (which is not affected by this dis-analogy) and found the priming effect.

[36] It should be noted that this experiment focused on comprehension, rather than production.

So: mental representations of linguistic elements form a network wherein there are stronger associations between some than others. Theoretical description of this has been developed in a number of spreading-activation models (particularly influential ones include Dell 1986; Levelt 1989; Roelofs 1992; Dell et al. 1997; Levelt et al. 1999), which theorize that activation spreads between mental representations based on associative connections; activating one representation (or "node" in the network) automatically leads to activating associated representations, and activation of one representation can be strengthened by activating other representations associated with it. These associations exist between representations of different words and between phonological and semantic representations of individual words. The basic idea is that when selecting a linguistic element, mental representations of many elements are active to varying degrees, and which element is selected depends on the relative activation levels of the elements. Most models describe the mechanism as involving purely excitatory mechanisms: factors like priming or accessing associated items increase the activation levels of mental representations, and these decay naturally over time. Another view (Oppenheim et al. 2010; also see Dell et al. 2014) is that both excitatory and inhibitory mechanisms are involved: some connections are strengthened (e.g. using a word in association with some conceptual material results in strengthening this mapping) and others are weakened (e.g. when a potential word is rejected; this means that a rejected competitor will be relatively less-strongly activated in subsequent language production).

4.2.3 Compartmentalization and the selection of elements

To relate all of this to the matter of the selection of elements: categorization is an important means by which elements are selected to express thought; the neural mechanism for categorizing a bit of conceptual material in the selection of elements may involve (in part) competition between activation levels of the mental representations of those elements. Now, the important part of this—on what basis the activation levels vary for a particular bit of conceptual material to be categorized—is still not understood. The phenomena discussed above relate to ways in which the

categorization of a bit of conceptual material can be biased, but not how it actually can be evaluated. If thought is restricted to fixed, reusable building blocks of conceptual material, the problem is much less complex, since individual building blocks can be associated with different elements. If, however, it is more flexible, and can contain 'novel' conceptual material—as intuitively appears to be the case[37]—, how the novel conceptual material gets associated with elements is a difficult question. A speculative (and not complete) answer is that the relation is established through mental model manipulation, as discussed in Chapter 3: this could allow contextualizing of the conceptual material in terms of possible linguistic elements and make their meanings be comparable, such that associations can be determined.

Now, while categorization is an important part of the selection of elements, it does not account for the entire issue. Categorization pre-supposes that there are bits of conceptual material that can correspond to linguistic elements and that there are linguistic elements that can correspond to bits of conceptual material (i.e., to use Langacker's terminology, that the thought is delineated into targets and that there are more-or-less suitable standards to match them). However, the available linguistic elements vary between languages and depend on the linguistic context within a language. For a given thought, different possible linguistic expressions divide up and organize the information to be expressed in varied ways; if the thought is delineated into targets, the available standards in the language cannot be depended on to match them closely enough that categorization can occur. There are often no elements that a(n uncompartmentalized) local bit of thought can be mapped to without invoking meaning construction and 'indirect' paths to meaning; meaning construction functions in a complex way that can yield related effective meanings from unrelated 'input' meanings (as seen in many of the examples above and in Chapter 3). If we consider, for example, that a particular idea can be expressed as either 'necessary' or 'not yielded', or either 'behold' or 'thoroughly hold', or either 'do A and B at the same time' or 'yībiān A

[37] It is conceivable, however, that at a low level, mental representations of apparently novel conceptual material are instantiated by combinations of reusable patterns of neural activity that do have associations with linguistic elements or similar low-level patterns comprising their mental representations.

yībiān B', it is apparent that categorization alone cannot possibly achieve the selection of the elements: the meanings of the elements in one expression and its alternative have no similarity; an account in terms of categorization would implausibly have the thought being categorized as either one thing or alternatively an entirely different thing.

My argument is that the process of the selection of linguistic elements for expressing a thought should be understood as a matter of compartmentalizing that thought, where the thought is divided up and organized according to what elements are available and how they construct meaning. Categorization is involved, since potential elements must be judged as being good or poor fits to varying degrees. But compartmentalization is what encapsulates the whole of the issue. Compartmentalization does allow the same thought to be expressed as either 'necessary' or 'not yielded': *necessary* or *not* and *yielded* and the construction that links them are available elements and they yield a meaning that is a suitable expression of the thought; knowing this allows the language user to divide up their thought appropriately.

4.3 Synopsis

In order to express a language-independent thought, a language user must compose an arrangement of linguistic elements that yields an approximation of that thought; part of this involves selecting the elements that are to comprise the arrangement, and part of it involves organizing the conceptual material in the thought so that those elements can be selected and put together.

Part of the cognitive mechanics of selecting elements involves the accessibility and activation levels of mental representations of the relevant elements, which can be influenced by frequency of use and exposure to use. The suitability of a particular element is determined partly by matching its meaning to a particular bit of conceptual structure, but what the bits of conceptual structure that are to be matched are depends very much on how the thought is organized in compartmentalization.

The account I have advanced here describes this organization of the conceptual material in the thought as being guided and constrained by what elements are available

and enabled by the language user's understanding of how the meanings of elements function and interact, which is what makes a particular organization/ compartmentalization able to approximate the thought.

To this end, there must be some cognitive device that makes possible the evaluation of the language-independent mental representation and the linguistically-shaped compartmentalized representation as being approximately equivalent. I think this device is related to mental model cognition.

While I do see compartmentalization as effecting the mapping of language-independent thought to linguistic elements, in a more generalized sense, it is simply the division and organization of conceptual information; it does not need to conclude with the identification of specific linguistic elements. In the context of creating an expression[38], for example, intermediate-level compartmentalizations (e.g. division into a phrase or clause), where the specific elements of the ultimate expression are not yet identified are possible. The outputs of these intermediate-level compartmentalizations are of forms such as 'some infinitive goes here' or 'a phrase to such-and-such effect goes here'.

While compartmentalization is describable as being comprised of two aspects—dividing up the conceptual material and organizing those divided-up bits—these aspects are often not separable. Also, because different parts of an expression interact with each other in cohering as part of an arrangement, compartmentalization has to be done holistically to an extent. However, this is not absolute, as I elaborate on in the next chapter.

Establishing the workings of compartmentalization as we did will enable us to describe the cognitive process of expressing thought in language in detail in Chapter 5.

[38] It is conceivable that the same cognitive operations are involved in discourse planning (i.e. dividing up and organizing some large-scale idea without reference (yet) to the specific elements that will express it) (here it may be comparable to subchunking), and perhaps even have some use outside the context of expression in language, but I will not speculate on this here.

Chapter 5

Abduction

Now we can directly address the central aspect of the cognition of expressing thought in language and the core of the theory presented in this dissertation.

We have identified that the expression of thought in language consists, fundamentally, of coming up with an arrangement of linguistic elements that will yield an approximation of the thought (c.f. Section 2.4). My central claim is that this is done through a kind of abduction that is driven by background knowledge of meaning construction. This is what we will discuss in this chapter. I set this up by giving the characterization that the understanding of the mechanics of meaning construction functions as a background theory on the basis of which the meaning yielded by a particular arrangement is predictable (in Chapter 3) and discussing how the linguistic elements that comprise the arrangement are determined through compartmentalization of the conceptual material that is to be expressed (in Chapter 4).

In Section 5.1, I will give a general introduction to abduction and distinguish abductive cognition (which is our object) from abduction as a logical operation (which is not quite it). Then, in Section 5.2, I will discuss the cognition of expressing thought in language. First, I will demonstrate that the fundamental cognition of expressing thought in language is a kind of abduction and specify the kind of abduction concerned. After this, I will discuss several important aspects of this cognition: mental model cognition, all-at-once versus step-by-step processing, the search mechanism, automatization, and unconscious processing and working memory. Finally, I will provide a model describing the abductive process of expressing thought in language in Section 5.3.

5.1 What is abduction?

5.1.1 General introduction to abduction

Very basically, *abduction* is the inference of a cause from an effect[1].

Abduction is one of the three basic kinds of inference, alongside *deduction* and *induction*. First, I will briefly contextualize abduction in comparison with deduction and induction.

Deduction is the "standard" form of inference; it consists of the inference of an effect from a cause. Peirce (1931-1935: 2.263) gives the example of concluding from the presupposition that all of the beans in a certain bag are white and knowledge that a certain collection of beans are from that bag that all of the beans in that collection are white. Other forms of deduction include the synthesis of the statements 'A is B' and 'B is C' into the statement 'A is C'. Mathematical proofs, for instance, are written in the form of deductions. Deductive problems can be given the following formalization:

What does p yield?

or

$p \rightarrow ?$

Induction is essentially pattern extraction. Peirce (ibid.: 2.619) gives the example of removing a handful (etc.) of beans from a bag, ascertaining that two-thirds of the beans in this handful are white, and supposing from that that two-thirds of *all* of the beans in the bag are white. A few clarifications regarding inexactness in what is meant by "induction" and how it is characterized are needed. One is that certain authors (e.g. Carnap 1962; Ladyman 2002: 28; Ormerod 2010) refer to all inference that is not

[1] The use of the term *abduction* for this sort of inference is due to Peirce (1931-1935).

In linguistics, some mechanisms of interpretation, language learning, and historical change have been described as abduction (see, among others, Andersen (1973), Deutscher (2002), Chomsky (2006), and Sakama and Inoue (2016)). Abduction in these contexts is used, essentially, for figuring out what particular expressions mean; this is quite distinct from the problem addressed in this dissertation, namely the expression of thought in language, which involves figuring out how to express a particular meaning. Additionally, unlike accounts of those issues, this dissertation is concerned with the cognitive mechanics of abduction.

deduction as "induction"[2]. I do not follow that convention primarily because, as we will see, what I mean by 'abduction' is not what I just described as 'induction'. Another is that induction is sometimes characterized as transferring past properties or regularities to the future (e.g. Schurz 2008: 202, following Hume). This characterization does not have sufficient generality: extrapolation from past to future is a special case of pattern extraction.

Abduction is the inference of a cause from an effect. It can be considered to be 'anti-deduction'. For example, solving the mathematical problem *What is 3 + 2?* uses deduction; on the other hand, solving *What equals 5?* uses abduction.[3] The basic form of problems in which abduction is used is the following:

What would yield q?

or

? → q

Syllogistic abduction can be schematized as $p \rightarrow q;\ q:\ p$ (i.e. the premises are that p yields q and that q is instantiated, and the conclusion is that p is instantiated), but, as we will see, this is unhelpful as a characterization of abduction in general. In the literature, there are a few different conceptions of what abduction is: in some views, its fundamental characteristic is identified as the production of an *explanation*, in others, it is described as the generation of a *hypothesis*, and in others, it is characterized as *inference to the best explanation*. Before clarifying what my own conception of abduction is (which I will do in Section 5.1.2), let us look at some examples of abduction to get a sense of how it works.

One example of abduction is medical diagnoses. A diagnosing physician is presented with a symptom (or a set of symptoms) and must come up with an affliction that would yield the symptom(s). The effect (q) that is the input of the abduction is the

[2] The basis for this is the notion that the fundamental distinction is between certain inferences (equated with deduction) and non-certain inferences (equated with everything that is not deduction); see Flach and Kakas (2000) for discussion. It should be noted that this distinction does not actually align with deduction and non-deduction, as I explain in Section 5.1.2.

[3] In this, abduction as the opposite of *deduction* ('leading from') should be termed *induction* ('leading to'); unfortunately, as we have just discussed, *induction* means something else.

symptom (or set of them), and the cause that is to be inferred (p) is the diagnosed affliction. Abduction is what is used to come up with the affliction.

A similar abductive problem is the following task. One is presented with a set of reference words (e.g. *pine*, *crab*, and *sauce*) and must come up with a word that would form a compound with all of them (e.g. *pine<u>apple</u>*, *crab-<u>apple</u>*, *<u>apple</u>sauce*) (this is known as a *compound remote associates* (CRA) *problem*; this example is from Kounios et al. 2006). In the medical diagnosis case, the abductive problem posed is *What would yield this symptom?*; here, it is *What would yield a compound with the reference words?*. The effect to be solved for (q) is the set of reference words, and the cause that provides a solution for that effect (p) is the compound-forming word.

Another kind of abduction is seen in the creation of scientific theories. The effect to be accounted for (p) is the observed phenomenon, and the cause that produces that effect (q) is the mechanism posited by the theory. The problem to be answered is *What would yield this phenomenon?*. For example, if a scientist is presented with the phenomenon of redshift in astronomical observations (the effect), they may account for it by positing that movement away from the point of observation (the cause); this would result in 'stretching' of the wavelengths and thereby in the observed phenomenon of redshift.

A similar abductive problem is posed by the following mystery:

> There is a room, barren except for a shelf on which a cat is sitting, with no windows; the shelf is too high for the cat to jump to, and there is water on the floor. How did the cat end up on the shelf?

We want the cat on the shelf—this is the effect that needs to be accounted for (q). How do we get there?—this is the cause that produces that effect (p). An answer to this is that there was an ice block on the floor, from which the cat jumped up to the shelf, and which has now melted into the water on the floor: this would yield the effect (q).

Abduction is also used in many manipulations in mathematical thinking. For example, if one has before them the task *Add up all the digits from 1 to 50*, they may realize that there is a different (and more efficient) way to do the calculation. Rather

than (i) performing the dull and arduous sum directly (1 + 2 + 3 + ... + 50), they can perform the calculation by (ii) adding 25 pairs that each sum to 50, plus the middle number of 25 that these pairs converge to ([0 + 50] + [1 + 49] + [2 + 48] + [3 + 47] + … + [24 + 26] + 25). This new way of structuring the problem is a different mathematical representation that is mathematically equivalent to the 'old' representation (i): it yields the same mathematical meaning. Here, the question is *What would yield the sum?*, and the answer provided by the abduction is the new representation (ii). The effect to be solved for is the sum, and the cause that would yield it is the new representation of the sum.

In all of these, the task is to come up with a solution that would yield some effect. Now, these examples do not involve exactly the same kind of abduction. Before we discuss differentiation between different kinds of abduction (which we will get to in Section 5.2.1), I will distinguish between abduction as a logical operation and abduction as a matter of cognition. This we will do in the following subsection.

5.1.2 Abduction as a cognitive operation versus abduction as a logical operation

First of all, abduction as a kind of cognition is, simply, a particular cognitive process (or class of processes) that is/are used to solve abductive problems (such as those we just saw). To be explicit, an abductive problem is one of the form *What yields q?*. What this sort of cognitive process must do to solve these is produce a certain solution *because* that solution would yield q. A diagnosis of some affliction is made because that affliction would yield the symptoms presented; *apple* is produced because it would yield a compound with *pine*, *crab*, and *sauce*; movement away is posited because it would yield redshift; a melted ice-block is imagined because it would yield the cat on the shelf; the new mathematical representation is created because it would yield the sum in question. This is very much circular: evaluation of what the solution would yield is what drives the generation of that solution in the first place.

This is not quite the same thing as what is discussed by philosophers. They are generally referring to the navigation of a certain kind of logical system, not the

cognitive mechanisms that carry it out. Examination of abduction from the point of view of philosophy has produced characterizations of it as explanation of phenomena, hypothesis generation, and inference to the best explanation. While abduction as a cognitive process is crucially involved in all of these activities, none of them are suitable as descriptions of the cognitive process.

The characterization of abduction as the production of explanations of phenomena only applies to some cases: while sometimes abduction amounts to producing an explanation, in other cases, it does not. In the five examples of abduction given earlier (medical diagnoses, CRA problems, theory of redshift, the cat mystery, and mathematical sum), three of them—diagnoses, redshift, and the cat mystery—involve abducing what amounts to an explanation (e.g. the patient has such-and-such symptom because they have such-and-such affliction; redshift occurs because movement away stretches the wavelengths; the cat managed to get on the shelf because it jumped up from a block of ice). But the other two (CRA problems and the mathematical sum)—while they are no less abduction—involve no such explanation. Finding a word that would form a compound with *pine* and *sauce* has nothing to do with explaining anything; neither does finding another way of doing the 1+ 2 + ... + 50 sum[4]. This is why I always use the wording *yield q*, rather than *explain q*. In cases where the solution produced by the abduction can be said to "explain q", what is actually produced by the cognitive process is something that *yields* q. Mechanistically, abduction as a cognitive process does not look for something that explains the effect; that the solution amounts to an explanation (if it does) is a secondary consequence of it yielding the effect. Something yielding something else can be directly computed, but it seems unlikely that something explaining something else can be: it is a secondary consideration.

A hypothesis, as I understand it, is a tentative 'claim' that something is the case. There are thus two aspects to hypotheticality: the suggestion that something is the case

[4] Of course, if, suppose, a student asks *Why is true that* $\sum_{i=1}^{n} i = \frac{n}{2}(n+1)$*?*, it would be perfectly reasonable to respond by 'explaining' *Because, if you think about it,* $\sum_{i=1}^{n} i$ *can be broken up into* $\frac{n}{2}$ *pairs that sum to* $n+1$ *(if n is even), or* $\frac{n+1}{2}$ *pairs that sum to* n *(if n is odd)*. But the structuring of the representation itself does not explain why the sum has the value that it does.

and uncertainty in that suggestion. Neither of these apply to abduction as a cognitive operation.

With regard to the latter, hypotheticality (as 'uncertainty') does not characterize abduction, because it can be a property of the products of deduction and induction as well. For induction, there is certainly uncertainty in how well or completely a pattern can be extracted from limited observation; to refer to Peirce's bean example, the inference from the fact that a handful of beans has a certain proportion of white beans that the entire bag of beans has that same proportion is just a hypothesis. For deduction, while it is true that reasoning from *(definitely) p yields q* and *(definitely) p* produces a conclusion of logical certainty (*definitely q*), if the system instead involves *maybe p yields q* or *maybe p* (or both), then the conclusion can be only a hypothetical *maybe q*. If one knows that maybe all the beans in a particular bag are white and maybe a particular handful of beans was taken from that bag, one can conclude only that maybe all of the beans in that handful are white. Conversely, abduction in the examples of the CRA problem and mathematical sum given above are necessarily correct: p definitely yields q.

Perhaps more importantly, whereas "philosophical" abduction ends with an assertion of *p*, mine ends with an assertion of *p would yield q* (to talk about it in philosophical terms, of course); this has nothing to do with whether p and q are true or false. The cognitive process that I am referring to as 'abduction' does not produce suggestions that something is the case: it only produces causes that would yield certain effects. It is not concerned with the truth of p or q; since there is no suggestion that p is the case, there is no production of a hypothesis. Barring errors in reasoning and uncertainties (like those above), the solution produced by this cognitive process is necessarily correct, insofar as that p would necessarily yield q. Any consideration of the truth of this p (if it even occurs—as we saw in the paragraph about explanation versus yielding, this is often irrelevant) is a secondary epistemic judgment and not part of the actual abduction. Adopting the result of an abduction as a hypothesis is a separate process.

In the domain of the philosophy of science, a common account of (some aspects of) the progress of scientific understanding regarding actual hypotheses refers to the

notion of *inference to the best explanation* (originally proposed by Harman 1965). In the enterprise of science, scientists come up with hypotheses/theories to explain phenomena. And in some cases, more than one hypothesis might be possible or come to mind. In these cases, it is possible to evaluate the hypotheses and decide if one is better than another. Some philosophers (e.g. Harman 1965; Barnes 1995; Ladyman et al. 1997; Lipton 2004; Thagard 2007; Schurz 2008) have taken this process that results in the determination of what the best explanation is and called it or equated it with 'abduction'. Others (e.g. Campos 2011; McAuliffe 2015; Yu and Zenker 2018) argue that only the generation of hypotheses should be called 'abduction', and their evaluation is a separate process. Now, the defining characteristic of abductive cognition is that the solution is produced *because* it yields some particular thing. In this sense, the evaluation of the "hypothesis"[5] is intrinsic to its generation: the only hypotheses that are generated are good ones that do yield the effect or phenomenon in question. But this sort of evaluation is different from the retrospective evaluation involved in judging hypotheses that have already been generated; we are concerned with generating the thing that might be adopted as a hypothesis. So the cognitive operation I refer to as abduction is not characterized as inference to the best explanation: it does not produce explanations, it does not produce hypotheses, and it does not include retrospective evaluation of hypotheses[6].

Abduction, as I refer to it hereafter, is a cognitive operation that produces a cause that would yield some desired effect: a potential p is produced as a solution because it is something that would yield q. Essentially, it is something like a search for cognitive objects and combinations of them that might yield the desired effect. In the following section, after explicating that expressing thought in language is carried out by a

[5] As explained above, abduction does not actually produce hypotheses, but rather things that may—in some contexts—be adopted as hypotheses.

[6] The process that I describe in Section 5.3 does include a mechanism for carrying out a sort of 'optimization' of the closeness of fit of the expression to the thought. However, this mechanism is, strictly speaking, not part of the cognitive operation of abduction, but rather of the larger process to which is applied. That larger process, though, does thus have some affinity to an "inference to the best explanation".

particular kind of abduction, I will discuss in greater detail the mechanics of that kind of abduction.

5.2 Abduction and the expression of thought

Now that it has been made clear what is basically meant by 'abduction', we can get to the point: the cognition responsible for expressing thought in language is abduction. First, in Section 5.2.1, I will demonstrate that the cognition of expressing thought in language is a kind of abduction and specify the basic properties of this particular kind of abduction. Following that, in Section 5.2.2, I will discuss aspects of the mechanics of the cognitive operation constituting this kind of abduction.

5.2.1 Abduction for the expression of thought in language

As just discussed, abduction is a cognitive operation whose input is some target to be yielded (q) and whose output is some thing that would yield that target (p). The essential form of the problem that it solves is 'What would yield q?':

$? \rightarrow q$

Recall from Section 2.4 that the overall process of expressing thought in language is essentially the deriving of a product representation (the linguistic expression) from a starting representation (the original thought), and we are trying to discover how the former is derived from the latter. The process starts with the language-independent thought (Λ_α) and derives from it a linguistic expression (Λ_β) that is approximately equal in meaning:

$\Lambda_\alpha \Rightarrow \Lambda_\beta$

$\Lambda_\alpha \approx \Lambda_\beta$

The linguistic expression is constituted by an arrangement of elements ($\rho_1 + \ldots + \rho_n$); this arrangement yields, according to the mechanics of meaning construction, an approximation of the original thought:

$\Lambda_\beta = \rho_1 + \ldots + \rho_n$

$\rho_1 + \ldots + \rho_n \rightarrow \Lambda_\alpha$

The task of expressing thought in language is to come up with an arrangement of elements that will yield such an approximation. The problem to be solved, then, is 'What would yield a meaning (approximately) equal to the original thought?':

$? \rightarrow \Lambda_\alpha$

This is a problem that is solved with abduction: the target to be yielded (q) is Λ_α, and the solution that yields it (p) is Λ_β. The abduction carries out something like a search for linguistic elements and combinations of them that would yield a suitable meaning. This is the very basic essence of my theory.

Now, there are different kinds of abduction. In what follows, I will specify the properties of the particular kind of abduction that is used to express thought. I will not attempt to give an overall typology of abduction; I will only mention features that make clear what I want to explain.

Firstly, abduction in the case of expressing thought in language is similar to that in the case where p = q. From the point of view of formal logic, the p = q case is of course trivial. However, when p and q are different representations of the same thing, the problem is very much not trivial cognitively. Earlier, we saw that the 1 to 50 sum can be restructured into a pair-counting calculation using abduction. The abduction has to find a different way of representing the mathematical meaning; it must construct a different representation that is to the same effect. Here, the starting representation is the direct sum, and the product representation derived from it using abduction is the pairwise sum. Expressing thought in language likewise begins with a starting representation (the original thought) and derives from it a product representation (the linguistic expression) that has an approximately-equal meaning using abduction.[7]

Another aspect is identified by Magnani (2009): he distinguishes between cases where the solution is chosen from a pre-established list (like in medical diagnoses),

[7] The syllogistic abduction in $p \rightarrow q;\ q:\ p$ is well-known as being fallacious (this fallacy is known as that of *affirming the consequent*): while it is possible that p is true (since it is not inconsistent with the premises), it is not guaranteed to be true (since it is not known whether p is the only way of getting q—if it is the only way, then q being true would require p being true as well, and it would then be guaranteed). So inferences of this form are generally not necessarily correct. However, in the case when p = q, there is no fallacy; the inference is necessarily logically correct.

called *selective abduction*, and cases where the solution is newly created (like in the invention of scientific theories), called *creative abduction*. Abduction in the case of the expression of thought is generally creative. Although the solutions (i.e. the linguistic expressions) are formed from elements taken from a finite list (i.e. all of the linguistic elements in the target language), the *arrangements* of the elements that constitute the solutions do not exist originally; rather, they are created by the abduction. However, it is conceivable that in some cases, the thought and its linguistic expression follow a habitual pattern (certain chunks of thought are re-used, and expressed using certain expressions); in these cases, since 'prefabricated' bits of thought are being mapped to 'prefabricated' expressions, the abduction is more selective—a habitually re-used expression is selected to match a habitually re-used bit of thought. If there is any novelty—any imperfection in replication—in either the thought or the expression, then it is not perfectly selective, however: there is some degree of creativity.

Equally important as whether the solution is taken from a pre-established list or is created is whether the fact that a (potential) p would yield q is pre-established or created by the abduction.

In some cases of abduction, it is pre-established; these are therefore relatively straightforward. For example, in syllogistic abduction ($p \rightarrow q$; q: p), the fact that p yields q is contained in the premise as an explicit rule. In these cases, the abduction is driven by the rule: the rule provides the basis for the production of the solution (c.f. Schurz 2008). A physician performing a diagnosis has in their head (before they perform the diagnosis) a list of potential ailments that would each yield a particular (set of) symptom(s); the fact that each ailment would yield whatever symptom(s) it would is pre-established (at least for "simple" diagnoses). Likewise, with CRA problems, where the task is to find a word that would form a compound with the given reference words, knowledge of the fact that a word that would form a compound with a reference word would in fact do so is pre-established. In these cases, the abduction is essentially a relatively straightforward search.

This is very much not the way expressing thought in language works. There is no list of rules that specify what each possible arrangement of linguistic elements would yield. Rather, the meaning of an arrangement is dynamically determined by knowledge

of how the meanings of individual elements function and interact (c.f. Chapter 3). This knowledge forms a background theory on the basis of which the arrangement (p) yields a particular meaning (q). This is essentially what Schurz (2008: 213-216) describes as *theoretical-model abduction*, where the theory—rather than a pre-established rule—drives the abduction. Schurz gives the example of the creation of a theory of buoyancy on the grounds that on any given bit of volume of water (etc.), there is an upward force exerted by the surrounding water, which is necessarily equal to the weight of that bit of water; if an object placed in the water weighs more than the volume of water displaced by it (i.e. that would otherwise occupy that space), it will sink, and if it does not, it will float. Here, the background theory driving the abduction of the mechanism of buoyancy is the understanding of gravitational force, contact forces, density, volume, and how they interact. This mechanism does yield buoyancy, but—of course—it is not known that it would until it is thought up. The kind of abduction used in our earlier examples of accounting for redshift and solving the cat mystery also functions this way. All creative abduction involves, in addition to the creation of the solution, the establishment of the fact that that solution would yield the desired effect: since there are no pre-established potential solutions, there can be no pre-established rules stating what those solutions would yield.

Since the point of abduction is that p would yield q, establishing that it does is crucial. In this kind of abduction, the 'search' is considerably more complicated, as there is no delimited set of potential ps from which to choose and no pre-established determination of what would be yielded by a potential p. In my own description, this sort of abduction involves the manipulation of elements in a mental model[8], since this allows the simultaneous construction of a p and determination of what it would yield (and specifically, construction of a p on *the basis* that it would yield a particular thing), which is what appears to be required. I will discuss this further in Section 5.2.2.1 below.

So, to summarize, the kind of abduction involved in expressing thought in language has the following properties. One is that the cause and the effect are

[8] Recall that by 'mental model', I mean the manipulable system rather than a diorama.

approximately equal to each other (i.e. $p \approx q$).[9] Another is that the solution is not taken from a predetermined list of potential solutions/causes. Also, there is no pre-established set of rules saying what would be yielded by a potential solution. Related to this last point, the kind of abduction we are concerned with involves what I refer to as mental model cognition, which is what informs what would be yielded by a particular cause/solution. In order for abduction (in any case) to be possible, what would be yielded by a potential solution must be predictable somehow. This is why I emphasized the importance of predictability of the meaning yielded by an arrangement in Chapter 3.

The cognition used in other arenas such as mathematical manipulation, computer programming, and music composition is partially analogous to that used in expressing thought. They all involve abduction that has these basic properties. In mathematical manipulation, mathematical elements are flexibly combined to express a mathematical meaning, but there are no rules that connect the overall mathematical expression to its meaning (this is impossible to begin with, since there can be no finite list of all the mathematical expressions)—rather, the only rules concern the function and interaction of the individual elements. In programming, the programmer starts out with an idea of what they want the program to do, and then they have to use whatever resources are available in the programming language to get it to do that. The particular operations that can be utilized in the programming language have no pre-established link to the particular effect (i.e. the function of the program) that the programmer wants to achieve. So the way it is done is by reconceptualizing the problem in terms of operations that can be used in the programming language (i.e. compartmentalization). They have knowledge of how these operations function and interact, and this is what enables them to come up with some arrangement of operations that achieve the desired effect. The same principles apply to composition: the composer has to figure out how to manipulate

[9] A potential objection to this characterization is that abduction is an ampliative process (i.e. it adds new information), and the expression of thought in language ought not to be. However, abduction in the $p = q$ case is non-ampliative (unless one takes changes in representation to introduce new information), and abduction in the $p \approx q$ case is minimally ampliative. On the other side, the expression of thought in language can only be done approximately, and it involves re-conceptualization (which does have a degree of ampliativity), so it, too, is minimally ampliative.

musical elements (pitch, timbre, rhythm, texture, variation, harmony, etc.) to express some musical meaning. Again, there are no pre-established rules linking the elements to particular meanings that the composer wants to express; the meaning is emergent from the arrangement of these elements in a rather complex way (which is far out of the scope of this dissertation). The same basic mental model cognition makes this meaning be predictable and enables the composer to create the musical expression.

5.2.2 Mechanics of the cognition

In this section, I will consider the details of the mechanics of this abductive cognition. I make proposals regarding the role of mental model cognition, all-at-once and step-by-step processing, the search mechanism, automatization, and unconscious processing and working memory. Some I assert more securely (mental model cognition and all-at-once and step-by-step processing); some are more speculative (the search mechanism, automatization, and unconscious processing and working memory).

5.2.2.1 Mental model cognition

Predictability of what would be yielded by a potential solution is at the essence of all abduction. This must be obtained even in cases where there is no pre-established rule specifying what q would be yielded by a potential p. In my proposal, this predictability derives from background knowledge of how the elements of the system work, and mental model cognition is what implements this knowledge in the actual abduction, in order to obtain a specific cause that will yield some effect in the absence of rules that say what would be yielded.

The importance of mental model cognition in abduction can be clearly grasped in the cases of our earlier examples of abducing the mechanism for redshift and solving the cat mystery. As mentioned above, in these cases, there is no pre-established rule determining specific effects that would be yielded by specific causes. Movement away from the point of observation and the cat having jumped up from what was an ice block and is now a puddle would indeed yield their respective targets (the phenomenon of

redshift and the cat on the shelf), but the fact that they would is not pre-established. The mental representation of the concept of a block of ice does not originally include properties like 'thing that a cat can stand on, that can melt later, and could provide sufficient height for a cat to jump to a shelf otherwise out of reach'. That these are true of a block of ice is established by the creative manipulation of a mental model. The manner of the manipulation of the mental model is governed by a background theory (in these cases, understanding of the relevant parts of the physical world), and on the basis of that theory, in the mental model, movement away yields redshift by stretching the wavelengths and a cat jumping up from what was previously an ice block and is now water on the floor would yield the scenario of the cat on the too-high shelf and water on the floor.

In the case of expressing thought in language, the background theory is the understanding of the mechanics of meaning construction: an arrangement of elements yields an approximation of the original thought ($\rho_1 + \ldots + \rho_n \rightarrow \Lambda_\alpha$) based on the understanding of the function of individual elements and how they interact (as discussed in Chapter 3). This background theory is what drives the abduction that creates the arrangement itself by making the meaning yielded by the arrangement be predictable.

Abduction in mathematics, programming, and music likewise utilize mental model cognition. Mathematical elements (numbers, variables, operators, etc.) are arranged to effect a particular mathematical meaning; program elements (i.e. operations) are arranged to effect a particular function for a program; musical elements are arranged to effect a particular musical meaning. The elements are manipulated as constituents of a mental model in such a way that they yield the desired effect.

Thagard (2007) gives a proposal for the mechanics of abduction that is partially similar to mine. It may be helpful for understanding what I mean by mental model cognition in abduction to explain it from a different angle, so I will borrow his explanation. His explanation is in terms of mental representations of causality within the system under consideration. He refers to this system as a "mechanism"; understanding how that system behaves is rendered by mental representations of what element causes what effect on what. For example, an electrician has an understanding of

the mechanism of a light switch and its circuitry: what is connected to what and what effects are caused by the interactions of the elements (and thus, if there is a failure in the system, the electrician is in a position to figure out exactly where the failure is). This is similar to what I am referring to as 'knowledge of how linguistic meanings function and interact'. However, the "mechanisms" that Thagard is considering are all 'static': for a particular system, all of the causal relationships and the specific effects produced are fixed. In language as a system for expressing meaning, this is not true. As just discussed, there are no pre-established rules that specify what meaning is yielded by a particular expression, and there are many different ways of getting at a particular idea using the conceptual resources available in languages, as seen in Chapters 3 and 4. So the sort of mental model cognition I am appealing to here subsumes the static causality Thagard discusses, but abduction for the expression of thought in language generally does not involve fixed causal relationships. As I remarked in Chapter 3, meaning construction in language is flexible and often based on reinterpretation. This cannot be achieved through fixed causal relationships. The way I envision it working, to use Thagard's way of talking, is that the "causality" inherent in the meaning of every linguistic element is extremely schematic: firstly, the effect that an element produces is not specific to a certain particular other element (unlike, e.g., a light switch or a bicycle). The abduction of a theory of redshift is quite unlike the abduction of what is broken in a faulty light switch: the latter is done on the assumption of specific causal interactions, and the former is done on the basis of more general principles that enable *novel* specific causal interactions (like that movement causes stretching of wavelengths) to be posited. In order for Thagard's description to be extended to capture creative abduction with pre-established causal linking (i.e. the 'rules' I have been referring to) between the cause and the effect, the causal connections cannot designate specific effects in interactions between specific elements (e.g. moving the switch closes or opens the circuit (one specific causal relationship), which enables or prevents the flow of electric current to the light-bulb (another specific causal relationship), which heats or stops heating the filament (another causal specific relationship), etc.). Instead, each element will have a schematic "causality" that determines potential causal relationships with other elements.

5.2.2.2 All-at-once versus step-by-step processing

Abduction is an intuitively opaque cognitive operation; when an abduction is performed, the solution tends to come all at once and suddenly. But abductive problems often can be solved either all at once or through multiple successive operations. This has been shown by research on *insight* (also known as the *Aha!* or *Eureka* experience), the psychological phenomenon of suddenly realizing the solution to a problem (often seeming obviously correct). I will be referring to insight in later subsections as well; first, I will briefly explain insight here. Insight is studied using various kinds of abductive problems, which are known as *insight problems*. Some examples are given below.

> A stranger approached a museum curator and offered him an ancient bronze coin. The coin had an authentic appearance and was marked with the date 544 B.C. The curator had happily made acquisitions from suspicious sources before, but this time he promptly called the police and had the stranger arrested. Why? (Metcalfe 1986)

> Transform this into a correct arithmetic statement by moving only one matchstick:
>
> 
>
> (Knoblich et al. 1999; depiction from Danek et al. 2016)

> A landscape gardener is given instructions to plant four special trees so that each one is exactly the same distance from each of the others. How would you arrange the trees? (Metcalfe 1986)

> CRA problems: e.g. *pine*, *crab*, *sauce* (Kounios et al. 2006); *tooth*, *potato*, *heart* (Subramaniam et al. 2009)

The original assumption in insight research was that solving insight problems is necessarily accompanied by an insight effect (see Bowden et al. 2005 and Danek et al.

2016 for brief reviews). The idea was that problems of this type have to be solved in one step: either one knows the answer or one does not. But insight problems can be solved either with or without insight (Jung-Beeman et al. 2004; Bowden et al. 2005; Fleck and Weisberg 2013; Danek et al. 2016; etc.). The reason for this is that—while a particular (single) operation of abduction does seem to be necessarily all-at-once—it is possible to break a problem down into multiple separate operations. For instance, a CRA problem can be solved in a single step (i.e. solving for all reference words at once), but it can also be solved incrementally, by abducing a solution for one reference word (or a subset of the reference words) at a time and then checking this candidate solution against the remaining reference words. For example, for the reference words *pine*, *crab*, and *sauce*, one might abduce from *pine* the candidate solution *cone* (which would form *pinecone*), but *cone* does not work for the other reference words (*crab* and *sauce*), so a new candidate must be found. For the reference words *tooth*, *heart*, and *potato*, one might abduce the candidate solution *ache*, which works for *tooth* (*toothache*) and *heart* (*heartache*), but not for *potato*, so a different candidate must be generated. If a candidate solution is successful, it will pass all of the checks, but in this scheme, the candidate solution is not actually generated in consideration of all of the reference words—it just happens to work for all of them. So one might, in consideration of *tooth* only, come up with the candidate solution *sweet* (which would form *sweet-tooth*), and then proceed to check it against *heart* and *potato*, where it passes the checks (as it would form *sweetheart* and *sweet potato*), but *sweet* is generated as a solution only for {*tooth*}, not for {*tooth*, *heart*, *potato*}. If these operations are performed successively, the correct solution is not realized suddenly, and the insight effect is absent or reduced. Similarly, the production of a solution to the matchstick problem shown above can be broken down into two constraint relaxations (realizations that an assumption one has (unconsciously) made should be discarded)—recognition that the = and – elements can be manipulated and that the = can be decomposed—and thus solved without insight (Danek et al. 2016). Finding an easy way to do the 1 to 50 sum can also be done across multiple steps: one may first obtain a structural schema of the answer without having gotten the full answer; e.g. one may abduce the fact that the sum can be restructured by rearranging the objects being summed into some easily-countable number of pairs that

each sum to the same value, but still need to work out the details (like what they sum to and how many pairs there are) separately. So the expression of thought in language likewise can, in principle, be carried out either all at once or step-by-step. This will be discussed in detail in Section 5.3.

5.2.2.3 The search mechanism

I described abduction as something like a search for cognitive objects and combinations of them that might yield the desired effect: when one is faced with an abductive problem, one starts out with the desired effect and must find some arrangement of cognitive objects that would yield it. However, the search mechanism is a bit of a mystery. Since Peirce, it has been recognized that abduction (at least creative abduction) cannot be done by a trial-and-error search through a potential solution space (unless, trivially, the potential solution space is delimited by exactly the same process that effects the abduction): it is not a guess-and-check process[10]. How the search manages to be successful is a puzzle.

In the kind of abduction we are concerned with, any potential p must be created by the abduction. The formation of a potential p—and thus abduction itself—is a creative process where the p is constructed from available elements. It can thus said to be a sort of 'directed creativity', where creativity is directed toward the realization of some determined goal (in this case, a p that would yield q).

My hypothesis is that potential bits of the arrangement are mentally represented and automatically plugged into the mental model, which determines the meaning they yield; if they go toward yielding the desired meaning, they are then activated. The arrangement can thus be constructed by applying the mental model to potential elements as a sort of filter. It should be noted that this—if it is actually correct—would

[10] Peirce observes (1931-1935: 5.591): "How is it that man ever came by any correct theories about nature?...You cannot say that it happened by chance, because the possible theories, if not strictly innumerable, at any rate exceed a trillion—or the third power of a million; and therefore the chances are too overwhelmingly against the single true theory in the twenty or thirty thousand years during which man has been a thinking animal, ever having come into any man's head."

account for only part of the search mechanism, since it does not explain how the combinations of elements are generated[11]. Research on insight supports this idea: it is hypothesized, based on EEG and fMRI data, that during insight solution of CRA problems, potential solutions are represented in the brain and attention shifted to them prior to solution (Jung-Beeman et al. 2004; Bowden et al. 2005).

I caution, however, that insight should not be confused with the cognition of expressing thought in language: they are not identical. Insight is a psychological effect—the subjective experience of suddenly becoming consciously aware of a solution. While the cognitive process underlying insight solution and the cognitive process of expressing thought in language probably do overlap significantly, there are two major differences between the two. One is that the cognitive process of expressing thought in language that I propose does not require the expression to be obtained all at once; insight effects, however, appear to require all-at-once solution. Another is that whereas insight requires consciously recognizing the solution, expressing thought in language is generally automatized, and usually does not involve explicitly realizing that one has obtained the solution (regardless of whether the expression was abduced all at once or step by step). Now, in some cases of expressing thought in language, insight effects can still occur, presumably. Sometimes, the process of expressing thought is very much not automatized, and people can spend quite a long time deliberating on how to express themselves. In these cases, insight effects would probably occur to some degree, because there would be a tendency toward being at least semi-consciously aware of the progress toward the solution, and because insight effects are associated with overcoming an impasse (which would be experienced in these cases).

5.2.2.4 Automatization

The question of how the process of expressing thought in language obtains automatization is another major problem. Expressing thought in language is generally

[11] Ellamil et al. (2012) suggest that the medial temporal lobe may be involved in creative generation and recombination of ideas. The mechanics of how this would be achieved are unknown, however.

unconscious, fast, and easy (not always, of course, though). Many abductive problems, like the mathematical sum and the cat mystery, stand in contrast to the expression of thought in this respect: their solution is generally neither fast nor easy. The solution rates of similar insight problems are typically quite low—around 10% in some cases (c.f. Danek et al. 2016)—; CRA problems are not so hard[12], but their solution is still much less automated than expressing thought in language. There is not an absolute distinction between these and expressing thought, because these harder abductive problems can be solved unconsciously, and in some cases (and by some individuals) can be solved instantly (though these cases are exceptional), but still, there is a definite qualitative difference either between these kinds of problems and the problem of expressing thought in language or in some aspects of the cognition used to solve them. So an explanandum is how or why expressing thought in language is automatized. I do not have an account of this issue that I am certain of, but I do have some speculative ideas that may partially explain the automatization of the process. I describe three proposals (which are mutually compatible) below.

The first involves handling a large amount of the burden of figuring out what would be yielded at the stage of accessing concepts from memory. If access of potential elements in the expression is done through a 'backward' lookup based on what they would yield, what they would yield would not have to be computed in its entirety at a post-access stage. This would work by having encyclopedic information pertaining to the function and interaction of concepts associated with the elements function as the basis of retrieval of those concepts from memory. For example, <outside> (as in *outside my specialty*) can be accessed on the basis that conceptually, if something is outside, it is separate and distinct. The filter selection mechanism described above would thus be able to automatically access the relevant elements even if they are not 'straightforwardly' related to the desired meaning.[13] Aronowitz (2019) proposes the

[12] In fact, one of the motivations for the introduction of CRA problems into insight research, by Bowden and Jung-Beeman (2003), was the high difficulty of the classic insight problems, which posed methodological problems due to the small number of data points that could be collected.

[13] Of course, it would then need to be explained how encyclopedic information is encoded and processed.

interesting idea that memory may function as a modeling system; if this is the case, it may facilitate this method of simultaneous access and 'computation' of meaning that I am describing. Aronowitz suggests that memory access could direct the processing of information: "my proposed change to the internal workings of memory also would lead to changes in the external role of memory, by which I mean the way other parts of our cognitive faculties are hooked up to memory. A rough way of characterizing this...shift is to think of how we use a set of facts from a file-box versus how we use a model. Where the former gets inputted and digested, the latter acts more like a guide." (p. 500). Unfortunately, however, she does not elaborate beyond this on how this would work mechanistically.

The second is that associative access can be used to build up an arrangement in an automatized fashion. Elements have associations with other elements based on statistical patterns in co-appearance (c.f. Janssen and Barber 2012); the idea here would be that accessing one element or a few elements will result in automatically accessing other elements and patterns of organization that are statistically associated with them. For example, if <bicycle> is accessed, <ride> will (tend to) be as well, due to the statistical tendency for expressions that use either *bicycle* or *ride* to use the other as well. In effect, this removes the need for all of the elements in an expression to be individually and separately accessed. This applies to some constructions as well (although many probably have no such associations). For example, the construction <VERB one's way to X> is quite strongly associated with *make* (as in *He <u>made</u> his way to the cake*). So accessing the construction as part of the abduced solution may tend to activate *make* as well, without having to find it separately. If associative access is responsible for part of the automatization of the process of expressing thought in language, when the association between the elements is stronger, the process of expressing thought would be more automatized, and when it is weaker, it would be less automatized. So, where *He <u>made</u> his way to the cake* has a strong association between the elements (here, the verb and construction), cases like *The tennis player <u>served</u> her way to the final*, which use a less-canonically associated verb, have a weaker association between the elements. In the former case, the generation of the expression may tend to be more automatized, and in the latter case, it may tend to be less

automatized. The notion that associative access can be used in the production of the solution is consistent (again) with ideas developed from research on insight. A hypothesis developed from EEG and fMRI data on the solution of CRA problems is that some potential solution words are more-strongly associated with the reference words and some are less-strongly associated with them; cognitive control (the management of cognitive processes) may be used for detecting solution candidates and using or shifting attention to (particularly less-dominant) associations within the brain (Kounios et al. 2006; Subramaniam et al. 2009; Kounios and Beeman 2014).

The third is that compartmentalization can (partially) follow learned patterns, so the thought does not always need to be given an 'original' or 'inventive' compartmentalization, but rather can use a 'prefabricated' one. In order to use a language fluently, one must be able to flexibly and fluidly compartmentalize one's thoughts in certain patterns; acquiring these patterns of compartmentalization is part of language acquisition, since different languages use different patterns. A possibility is that this ability to flexibly and fluidly compartmentalize thoughts in certain ways is acquired partly by extracting patterns found in compartmentalizing various thoughts in various ways in day-to-day use of the language. This could involve certain recurrent features or aspects in the thought being associated with certain linguistic elements or patterns of structuring them, and thus could lead to a degree of (schematic) prefabrication in the construction of arrangements of elements.[14] This idea of the importance of associations with extracted patterns is consistent with the conception of linguistic knowledge advanced by Taylor (2012); he emphasizes linguistic knowledge as patterns extracted from language use. I will provide a detailed description of how compartmentalization could be integrated into the abduction of the expression in Section 5.3.

[14] A related, but distinct possibility is that as an extreme case, the language user's patterns of thinking themselves would adapt to the target language (c.f. Levelt 1989; also Slobin 1996). I will not rule out that this might occur in some cases, but I am assuming the general case where thought is independent of language.

5.2.2.5 Unconscious processing and working memory

Now, it may be important that the process of expressing thought in language is generally unconscious. The reason has to do with working memory capacity. Working memory is a temporary store of information for the purposes of manipulation, which has a severely restricted capacity, generally somewhere on the order of 4 items (Baddeley 2003, 2007, 2012; Cowan 2008; Eriksson et al. 2015). In the sort of abduction we are interested in, a large number of potential interactions between a large number of potential elements must be considered, which may vastly exceed this limit. But unconscious processing may not be so constrained. People do presumably think unconsciously about several things at once. Relatedly, introspective accounts of work on very hard abductive problems, like in mathematics, commonly describe an *incubation* period (so termed by Wallas (1926)) of unconscious rumination, where (presumably) elements that are to constitute the solution are unconsciously manipulated[15] (e.g. Poincaré 1910; Hadamard 1945; Koestler 1964, particularly Ch. V; c.f. Ghiselin 1952; Smith et al. (eds.) 1995; Csikszentmihalyi 1996); the technique of incubation has been shown to be beneficial for creative problem-solving (see Sio and Ormerod 2009 for an overview), and insight is "preceded by substantial unconscious processing" (Kounios and Beeman 2014: 88). Similar conjectures about unconscious processing not being so constrained in capacity of manipulated items have been made by Hassin et al. (2009), based on the fact that "there are points in time in which we

[15] This is not necessarily the case, however. Other suggestions for what is going on in incubation are that it is simply recuperation from fatigue, that conscious work on the problem is performed intermittently during the incubation period, and that leaving the problem alone can weaken *mental sets* (these are a generalized case of the constraint relaxation mentioned earlier in Section 5.2.2.1) that were not productive for solving the problem and allow re-approaching the problem from a new perspective. See Gilhooly (2016) for a review; also c.f. Hadamard (1945: 32-37), Smith (1995), and Allen and Thomas (2011). Gilhooly argues for the hypothesis that incubation is a stage of unconscious work on the problem; Yuan and Shen (2016) and Shen et al. (2018) argue for a mixture of conscious and unconscious processes (in Shen et al.'s words, "insightful incubation and even some incubations in other modes of creative thought such as divergent thinking may depend on some yet to be determined interaction between conscious cognitive control and unconscious processes" (p. 197)).

seem to be advancing multiple goals, decisions and plans (etc.)" (p. 667), and by Dijksterhuis and Nordgren (2006), who report experimental evidence suggesting that complex problems are better solved unconsciously; Hadamard (ibid.) also proposes that unconscious processing is necessary for the manipulation of large numbers of elements[16]. Dijksterhuis and Meurs (2006) have also pointed out that another potential advantage of unconscious processing is that it is less restricted and more flexible in the space of potential elements to be considered; conscious thought, on the other hand, has a tendency to fixate.

One possibility for how unconscious processing may be related to avoiding limitations in working memory capacity is that there is such a thing as unconscious working memory, characterized by both the content and the manipulation of the content being unconscious, that is not so constrained in capacity. On the whole, work into whether unconscious working memory exists is inconclusive, particularly for the sort of system I am appealing to, which requires unconscious manipulation.

Working memory has traditionally been conceived of as conscious (e.g. Baars and Franklin 2003; Baddeley 2003, 2007). However, recently, there have been several investigations into the possibility of unconscious working memory. A number of studies have shown evidence for the storage and short-term maintenance of unconsciously-perceived information (Hassin et al. 2009; Soto et al. 2011; Bergström and Eriksson

[16] Hadamard remarks: "[W]e see that the unconscious has the important property of being manifold; several and probably many things can and do occur in it simultaneously. This contrasts with the conscious ego which is unique. We also see that this multiplicity of the unconscious enables it to carry out a work of synthesis...it is obvious that invention or discovery, be it in mathematics or anywhere else, takes place by combining ideas. Now, there is an extremely great number of such combinations, most of which are devoid of interest, while, on the contrary, very few of them can be fruitful...to find these, it has been necessary to construct the very numerous possible combinations, among which the useful ones are to be found. It cannot be avoided that this first operation take place, to a certain extent, at random, so that the role of chance is hardly doubtful in this first step of the mental process. But we see that that intervention of chance occurs inside the unconscious: for most of these combinations—more exactly, all those which are useless—remain unknown to us. Moreover, this shows us again the manifold character of the unconscious, which is necessary to construct those numerous combinations and to compare them with each other." (pp. 23-30).

2014, 2018; Trübutschek et al. 2017; Trübutschek et al. 2019a; also see Soto and Silvanto 2014 for an overview). However, while these results are consistent with an unconscious working memory, it is not decisively clear whether they actually involve 'true' unconscious working memory, and not unconscious short-term memory (this point is also acknowledged by Bergström and Eriksson 2014; see Persuh et al. 2018 for a review; also c.f. Eriksson et al. 2015: 43). The manipulation of the information is crucial. Trübutschek et al. (2019b) investigated the manipulation of unconsciously perceived information, and concluded that while the storage of information in working memory does not require consciousness, the manipulation of working memory content is associated with conscious processing (because they observed a neural signature of conscious processing coincident with the manipulation). However, their experimental paradigm involved *deliberate* manipulation[17], so it is not particularly surprising that a signature of conscious processing should be observed. This is different from non-deliberate processing. The unconscious processing I am suggesting is not deliberate: it is unconscious manipulation of information, not conscious manipulation of unconscious information. So this result does not necessarily apply to the sort of processing we are concerned with. Inaccessibility to awareness is insufficient; whether the process is driven consciously or not is also a relevant factor. At any rate, it does seem that unconscious problem-solving does occur, and it also appears that problems whose

[17] They briefly flashed a square in one of 24 positions along an imaginary circle on a screen, then showed a mask (an overlay that 'hides' the location of the target square), and directed participants to mentally rotate the square either 120° clockwise or 120° counterclockwise. On some trials, the square was dark, and on others, it was faint and barely visible. Participants reported whether they perceived the square or not. Participants were directed to guess at where the location of the rotated square would be if they could not see the initial target square; they guessed the location better than chance, in line with the results from the studies mentioned above showing that that non-consciously-perceived information can be stored temporarily. The authors also looked at a neural signal of conscious processing (desynchronization in alpha and beta frequencies; this desynchronization distinguished between seen and unseen trials) after the presentation of the direction to rotate clockwise or counterclockwise, and found that it was present on both seen and unseen trials. From this, the authors inferred that the information was transferred to conscious working memory for the manipulation.

solution would exceed working memory capacity are indeed nevertheless solved, which suggests that in some situations, the ~ 4-item limit is not a constraint.

Another possibility is that the capacity limits of (conscious) working memory arise from limitations in the focal scope of attention (Cowan 2008; also see Dehaene and Naccache 2001; Dehaene et al. 2006; Kouider and Dehaene 2007). The idea here would be that operations on long-term memory items and working memory items are essentially the same. There are myriad manipulations being conducted at any one time, and the items in working memory are simply the subset of the active long-term memory items that are being attended to (Öztekin et al. (2010) also suggest that there is a single store for long-term memory and working memory representations). There is no distinct 'working memory' store for either conscious or unconscious processing. Here, what is called 'working memory' is conscious by definition, but manipulations can also be done on the active portions of long-term memory (which would be unconscious); manipulation is not the exclusive privilege of working memory. A crucial point here is that the active long-term memory is not so limited in capacity. As a result, non-conscious processing would not be constrained.

So we have two alternative conjectures and a null hypothesis:

i) Unconscious working memory capacity is significantly higher than conscious working memory capacity, and so it is important that the conversion of thought to language is generally unconscious.

ii) It is not particularly 'important' that conversion of thought to language be unconscious (since the same manipulations can be done on consciously-attended objects and unconscious objects), but generally is so because the potential elements being manipulated cannot fit into the scope of attention.

iii) Unconscious processing has no significant consequences for the relevant cognitive processes, and so it is not particularly important that the conversion of thought to language is generally unconscious (perhaps it is nothing more than a result of the conversion process being fast and easy).

5.3 A model of the abductive process of the expression of thought in language

In this section, I will describe a model of how the cognitive operation of abduction can be used to derive the linguistic expression from the original thought. The more speculative aspects of the depiction of the cognition discussed above (i.e. the search mechanism, automatization mechanisms, and the role—or lack thereof—of working memory) are compatible with this model, but the model does not depend on them.

As discussed above, abductive problems can be solved either all at once or through multiple steps. There are thus two paradigms that I propose: the *all-at-once paradigm*, and the *algorithmic paradigm*. In the all-at-once paradigm, the entirety of the expression is abduced at once; in the algorithmic paradigm, the expression is abduced bit by bit.

As we saw in Section 4.1.2.2, when creating a linguistic expression, there are multiple requirements that must be simultaneously satisfied. This property is not exclusive to language; it is also well-illustrated, for instance, in musical composition. A composer writing a contrapuntal passage must do so in such a way that each musical line says what it ought to, that the harmonies resulting from the overlay of the various lines say what they ought to, and that the overall effect emergent from the composite structure is what it ought to be. These layers of musical structure (the individual lines, the harmonies, and the overall emergent effect) are interdependent and must be balanced against each other during the compositional process, in a way similar to that in which interdependent bits of linguistic expressions must be taken together to construct a linguistic expression. This property of having various interdependencies that must be simultaneously satisfied—whether in music or language or whatever—is partially structurally analogous to the property of CRA problems that the three (or however many) items must be simultaneously satisfied. The two paradigms I propose represent the two essential ways that this can be handled.

The algorithmic paradigm is analogous to the multiple-successive-abduction solution method. For CRA problems, searching for solutions for individual reference

words produces candidate solutions that are then checked against the remaining reference words; here, candidate expressions for local bits of the thought are abduced individually and successively and then checked for compatibility with each other. The discussion below will focus on elaborating the details of this paradigm. My suggestion for the all-at-once paradigm is that it uses the same cognitive processes as the algorithmic paradigm, but performed in parallel instead of in succession. The CRA analog of this solution method would be performed by carrying out parallel searches for compounds of each of the three (or however many) reference items, with potential solutions not activated unless they match across all of the parallel operations.

The basic idea is that as a first step, a tentative compartmentalization of part of the thought is made. This produces a *candidate expression*, which serves as a pivot from which successive bits of the expression are generated. Once a candidate expression is tentatively set, the next candidate expression is abduced; this continues until the entire expression is constructed. The abduction of each successive candidate expression is made on the assumption that it will be integrated into the all of the already-produced tentatively-set candidate expressions.[18]

To give a very simple example, the initial compartmentalization of the thought is made, and [happy] is chosen as a candidate expression that would yield that local bit of meaning. [happy] by itself, however, is inadequate as an expression of the entire thought, so the algorithm repeats: *[happy] in combination with what?*. As a result of this, another candidate expression is produced, and so the tentative expression is now [happy and energetic]. This is satisfactory as an expression of the thought, and so the process is completed.

[18] It should be noted that, if working memory is involved in the expression of thought in language, the algorithmic paradigm would be significantly less demanding on working memory, as only the tentatively-set part of the expression and potential elements that may constitute the candidate expression under immediate consideration would need to be placed in working memory; the remaining candidate expressions yet to be addressed are not under consideration, and so information relating to them does not need to be placed in working memory. The information that is in working memory, however, may still be substantial, involving a large number of potential elements, so the algorithmic paradigm is not an automatic solution to the limited-working-memory-capacity problem.

The basic process of the algorithmic paradigm is sketched in Figure 9 below (c.f. Figure 6 in Section 2.4). Each iteration of the algorithm produces a different candidate expression that forms part of the eventual linguistic expression. The original thought presets an abductive problem that is to be solved. A candidate expression is abduced for part of the thought[19] and is tentatively set. The problem to be solved is then updated to include the tentatively-set candidate expression. The algorithm then repeats: a new candidate expression is abduced for this updated problem. This new candidate expression is then added to the tentative expression. This process repeats until the entire expression is completed.

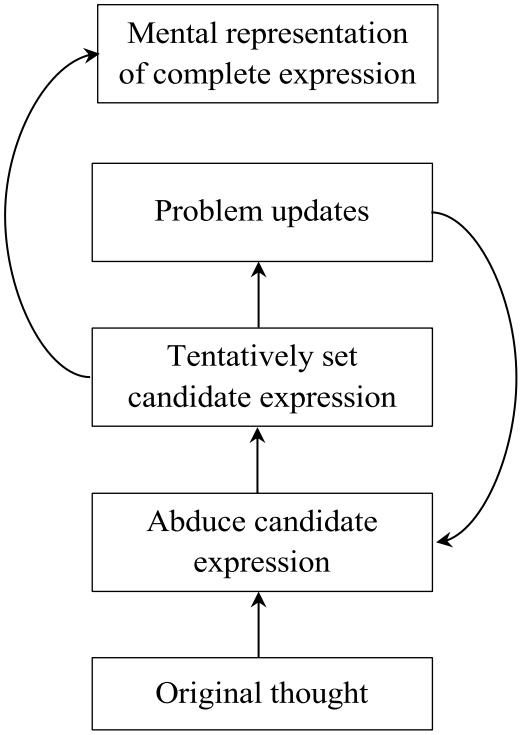

Figure 9: The basic outline of the algorithmic paradigm

Let us illustrate how this would work with a specific example. The order of the candidate expressions used in this following illustration is arbitrary; in principle, the

[19] Compartmentalization relates the solving of the abductive problem to the possible solution space. Every time a new candidate expression is to be abduced, the conceptual material corresponding to that candidate expression has to be compartmentalized.

candidate expressions that are to form the linguistic expression can be abduced in any logically possible order. I discuss the issue of order in more detail later a few paragraphs below.

The derivation of the expression *some people* could be accomplished as follows. The language user abduces [some] as the first candidate expression. This is tentatively set as part of the expression being constructed, and the abductive problem updates to *What in combination with [some] would yield (an approximation of) the original thought?*. The language user must then solve this new problem through a second operation of abduction. The second abduction produces the <QUANTIFIER NOUN> construction [20] as the next candidate expression and this candidate expression is tentatively set[21]. The problem updates again. A third abduction produces [people] as the

[20] As discussed earlier (Section 2.2.1.3), the generality of constructions often depends on the individual's analysis. What I am referring to here as the <QUANTIFIER NOUN> construction may not actually be analyzed as such: the actual analysis of the language user may have the construction in question be specific to *some*, for instance (and distinct from constructions involving other quantifiers, like *many* and *few*).

[21] In principle, if one or more candidate expressions have been tentatively set, a new candidate expression can be abduced either as i) a direct solution to the updated problem as a whole (that is: the calculation takes into account the tentatively-set candidate expression(s)), ii) only as a solution to the local corresponding bit of thought in question (that is: the calculation does not take into account the tentatively-set candidate expression(s)), or iii) (if there are multiple tentatively-set candidate expressions) as a solution to part of the updated problem (that is: the calculation takes into account only some candidate expressions).

In each of these cases, the way the new candidate expression is integrated with (i.e. 'plugged in' to) (c.f. Langacker 2008: Ch. 6 (pp. 161-182), 2009: 68-73) the tentatively-set candidate expression(s) is different.

In the first case, the new candidate expression is from its conception integrated with the entire tentatively-set arrangement of candidate expressions. So in our current example, the <QUANTIFIER NOUN> construction produced as a solution would assume that the previously-obtained [some] is the argument in the <QUANTIFIER> slot.

next candidate expression. This too is tentatively set. At this point, the expression is completed, and the process ends.

The compartmentalization of the same conceptual material can be different, resulting in a different set of candidate expressions and updated abductive problems. In Japanese, for instance, the equivalent expression may be *...hito mo iru* (lit. 'there are also people who...'; *hito* means 'person'; *mo* is an additive particle; *iru* can be glossed as 'be'). In this case, instead of <some X>, the <X mo iru> construction is used. So the first candidate expression could be this construction. (For the purposes of this illustration, I will assume that this construction consists of the syntactic pattern alone[22].)

In the second case, since the new candidate expression is abduced as a separate ingredient, without consideration of the tentatively-set candidate expression(s), it is initially *not* integrated with them. It must instead be integrated when it itself is tentatively set. This amounts to an indirect solution to the updated problem, as whether it succeeds as a solution to it must be checked in a separate step.

In the third, intermediate case, the new candidate expression is initially integrated with some of the previously-obtained candidate expressions, but not integrated with others. Full integration is achieved when the new candidate expression is tentatively set (just like the second case above).

There are differences in the way these procedures fail. The second and third cases can produce candidate expressions that are incompatible with others, whereas failure in the first case can only result in no candidate expression being found as a solution. There is no reason to rule out any of these as possibilities at present. The attestation of certain kinds of production errors, which involve incompatible elements, strongly suggests that the second and/or third exist, at least. Errors and how they can be overcome will be discussed below.

[22] As I discussed in Section 2.2.1.2, it is possible in some cases for constructions to be associated with overt elements. In that section, I gave the example of the <to X> construction (as in *It's not a big deal to me*) being coupled with the overt element [to]. With regard to the current example, the construction involved in the expression *some people* may be analyzed as being coupled with [some]. In this case, the construction and [some] would not be abduced as separate candidate expressions; they would be accessed together. Similarly, it is possible for the syntactic pattern in <X mo iru> to count as an element alone (distinct from any overt elements), in which case the syntactic pattern and the overt elements ([mo] and [iru]) would be abduced separately as different candidate expressions (but see footnote 23 below). But it is also possible for the syntactic pattern and [mo] and [iru] (or perhaps just [mo]) to be coupled in the language user's analysis and abduced as a single candidate expression.

This is tentatively set and the problem updates. The language user abduces the next candidate expression, producing [mo]. This is tentatively set and the problem updates again. The next abduction produces [hito]. Again, this is tentatively set and the problem updates. The last abduction produces [iru]; this is tentatively set. With this, the expression is completed.

Candidate expressions can be lexical (i.e. words or morphemes) or constructional (i.e. syntactic patterns); as discussed in Chapter 2, there is no principled distinction between these types of elements, as they are both tools used to construct meaning, and there is no primacy between them. This is consistent with Bock and Ferreira's (2014) account that language production is both *word-driven* (i.e. the expression is built around words) and *structure-driven* (i.e. the expression is built around syntactic patterns).

There is no fixed order in which the expression must be built up, but the order is not necessarily random. There is psycholinguistic evidence suggesting that in some cases, cognitive salience can influence what is put first in a sentence (cued items tended to be made the subject and the constructions used tended to alternate between active and passive depending on whether the cued item was an agent or patient), though this effect was not seen in other cases (see Konopka and Brown-Schmidt (2014) and Bock and Ferreira (2014) for summaries). However, it is unclear if this reflects the temporal order of linguistic encoding (i.e. it is the subject because it is done first) or is a consequence of linguistic structure (i.e. it is done first because it is the subject); these possibilities are confounded by the subject-first structure of English. Also, following the syntax of the target language is beneficial for fluent production. For example, in the production of *The robber stabbed the man (with a knife)* (Brown and Dell 1987; also see Konopka and Brown-Schmidt 2014: 9-10), it is possible to add the information about what instrument the robber used to stab the man after formulating the 'core' of

the clause (*The robber stabbed the man*)[23], but this is not true of Japanese, where it must be said earlier (*dorobo ga naifu de otoko o tsuita*; word-for-word '(the) thief (a) knife with (the) man stabbed'). So leaving it for later will work for English, but not Japanese: the language user will get stuck if they have not figured out how to express 'with a knife' in time and may exhibit a dysfluency (i.e. a pause, a filler like *uh*, or a rewording).

We discussed earlier (Section 5.2.1) that expressing thought in language is essentially solving an abductive problem, $? \rightarrow \Lambda_\alpha$ (i.e. *What would yield (an approximation of) the original thought?*). A fundamentally important point to note is that in the algorithmic paradigm, this problem is not static; rather, it evolves as successive candidate expressions are produced and tentatively set. After the first candidate expression, ρ_1, is produced and tentatively set, the abductive problem becomes $\rho_1 + ? \rightarrow \Lambda_\alpha$ (i.e. *What in combination with ρ_1 would yield Λ_α?*); after the second candidate expression is produced and tentatively set, it becomes $\rho_1 + \rho_2 ? \rightarrow \Lambda_\alpha$ (i.e. *What in combination with ρ_1 and ρ_2 would yield Λ_α?*); this updating of the problem continues throughout iterations of the algorithm until the expression is completed[24]. In general, when there is an already-decided part of the expression (for example, when the language user has begun speaking before having decided the entire expression, or is

[23] Brown and Dell's analysis is that if the specific instrument is conceptualized as a basic part of the main idea, it is linguistically encoded in the same clause; if, instead, it is an "embellishment", it goes in a separate clause (e.g. *The robber stabbed the man. He used a knife.*). My point here is specifically about the temporal order: even if the conceptualized semantic role of the instrument is such that it goes in the main clause, it does not have to be determined (either in the thought or in terms of how it is to be encoded linguistically) until after the core of the clause is produced.

[24] In a 'pure' algorithmic paradigm, candidate expressions are single elements. However, I do not think candidate expressions in general are restricted to single elements, as I explain at the end of this section. Now, I actually intend for this description of the updating of the abductive problem to apply in general (whether the candidate expression is a single element or not). I have, however, thereby perpetrated a bit of abuse of notation here. Earlier (Sections 2.4 and 5.2.1), we gave a schematization of an expression as an arrangement of elements that yields an approximation of the thought ($\rho_1 + \ldots + \rho_n \rightarrow \Lambda_\alpha$), where individual elements are denoted by ρ_1, ρ_2, etc.; here, I am using ρ_1, ρ_2, etc. to denote individual candidate expressions instead. So here, what I refer to as ρ_1 (for instance) is not necessarily a single element.

speaking while they develop the thought to be expressed), the language user must build the remainder of the expression around that part (or else withdraw that part and rephrase); so, if we denote this already-decided part as ρ, the problem to be solved in these cases is $\rho + ? \rightarrow \Lambda_{\alpha}$. Anderson et al. (2009) propose that the anterior cingulate cortex is involved in setting subgoals in the solution of insight problems; if this is correct, it may also be involved in this updating of the abductive problem in the expression of thought.

Each candidate expression constrains the compartmentalization for all other candidate expressions that may interact with it. For example, if [can] is chosen as a candidate expression for a particular bit of thought, [one X imagine] and [it X be imagined] are possible as other candidate expressions, but [it is X to imagine] is not; conversely, if [possible] were chosen instead of [can], [it is X to imagine] would be possible and [one X imagine] and [it X be imagined] impossible. So every time the abductive problem updates, the possible space of subsequent candidate expressions that interact with tentatively-set ones is constrained. In Section 4.3, I briefly referred to 'intermediate-level compartmentalizations', which do not result in specific linguistic elements being designated, but rather relatively vague determinations like 'some infinitive goes here' or 'a phrase to such-and-such effect goes here'. Tentatively-set candidate expressions result in intermediate-level compartmentalizations of this sort for any subsequent candidate expressions that interact with them.

I have referred to candidate expressions being 'tentatively set' before subsequent candidate expressions are produced. The 'tentative' part is important, because in the progress of algorithmic generation of an expression, incompatible interdependencies do arise (some candidate expressions may not interact properly; c.f. Section 4.1.2.2 (Interdependencies in compartmental structure)), just as candidate solutions in CRA problems that work for one reference word may not work for another. In these cases, any candidate expression can be discarded. An example of this is shown below (for an abortive attempt at *This is totally outside my specialty*). Here, for clarity of notation and simplicity of explanation, I have written the candidate expressions as lexical items. However, candidate expressions also include constructions; also, they do not exist in isolation—they are incorporated into the tentative arrangement. The interactions they

have with other candidate expressions in this tentative arrangement are the basis for the determination of compatibility and incompatibility.

ρ_1	my	OK
ρ_2	specialty	OK
ρ_3	at all	OK
ρ_4	outside	incompatible

When an incompatibility is found, one or more of the candidate expressions must be discarded. In this example, ρ_3 and ρ_4 are incompatible, so one of them must be discarded and the abduction redone. If we discard ρ_3, the abductive problem would then update to $\rho_1 + \rho_2 + \rho_4 + ? \rightarrow \Lambda_\alpha$; if we discard ρ_4, it would update to $\rho_1 + \rho_2 + \rho_3 + ? \rightarrow \Lambda_\alpha$.

Now, in some cases, an incompatibility will cascade. Suppose we have a tentative expression built from $\rho_1 + \rho_2 + \rho_3 + \rho_4 + \rho_5$, and ρ_5 and ρ_2 are incompatible, so we discard ρ_5; however, without ρ_5, there is no available linguistic element that can be combined with ρ_4 to yield the desired meaning. As a result, ρ_4 must also be discarded.

To give a concrete example, we saw earlier in Chapter 4 that unlike English, Japanese cannot express 'A is special in a way that B is not' in an analogous manner to 'A is special in the same way B is'. This is reproduced below.

i) a. A is special in the (same) way B is. (= (54a) in Chapter 4)

b. A wa B no you ni tokubetsu da. (= (54b) in Chapter 4)
A TOP B GEN like ADV special COP
(lit. 'A is special like B')

ii) a. A is special in a way that B is not. (= (55a), (66a) in Chapter 4)

b. A wa B no [¬ you] ni tokubetsu da. (= (55b), (66b) in Chapter 4)
A TOP B GEN NEG.like ADV special COP
(lit. 'A is special unlike B')

c. A no tokubetsusa wa B ni nai. (= (55c), (66c) in Chapter 4)
A GEN special.NMLZ TOP B in be.not
(lit. 'A's specialness is not in B.')

But suppose the compartmentalization in the algorithmic paradigm proceeds along this (English-like) course. It would then end up with the non-existent candidate expression [¯ you]; that is: there is no available element that can be incorporated into the tentatively-set expression to yield (an approximation of) the desired meaning. So, naturally, [¯ you] will be discarded. Since [you] has been discarded, the <B no you ni> construction ('like B') must be discarded as well. Now, this construction connects *B* to *tokubetsu* ('special'). In order to express a suitable meaning, there must be some way to connect *B* to *tokubetsu*; but now that it has been discarded, there is no way to do that. So eventually, the predicate adjective construction (<A wa ADJ da>) has to be discarded, and the entire idea re-compartmentalized (into something like *A no tokubetsusa wa B ni nai*; lit 'A's specialness is not in B'). Re-adjustment of the compartmental structure can involve shifting, separation/compression, or fundamental formulation, described in Section 4.1.2.1 (Basic patterns of variation in compartmental structure).

Discarded candidate expressions can be held in reserve, so some that were previously discarded may be reactivated following a reformulation. The determination of what to discard and what to keep can be made on the basis of how important the candidate expression's contribution of meaning is, how important the corresponding part of the thought is, how well it approximates the original thought, and how difficult it would be to reformulate the expression were it to be discarded. Some candidate expressions are more flexible in accommodating other candidate expressions than others; this means that incompatibilities are less likely to arise (c.f. Bock and Ferreira 2014: 32-35). For example, [give] accommodates both the constructions <give OBJECT to NP> and <give NP OBJECT>, whereas [donate] accommodates only <donate OBJECT to NP>: both *give some money to the church* and *give the church some money* are possible, but only *donate some money to the church* is; **donate the church some money* does not work (Ferreira 1996; c.f. Bock and Ferreira 2014: 34). Ferreira found that the production of expressions using more flexible linguistic structures was performed faster and with fewer errors.

This system of updating and discarding can also be used to fine-tune the expression to better approximate the original thought (as I mentioned in footnote 6). The initially-produced candidate expressions are not necessarily those that best

approximate the thought (it is possible that the first candidate expressions are biased toward those with higher accessibility; c.f. Section 4.2.2). Successive searches for more precisely-approximating candidate expressions can be done to provide a sort of optimization of the closeness of fit to the thought; alternative candidate expressions and tentative arrangements of them can be compared and the best-fitting one selected.

One type of failure in the process of expressing thought can be seen in cases where incompatibilities are not resolved, as in the errors shown below known as *blends* (these examples are from Bock and Ferreira 2014: 35). Each of the malformed expressions (i-iii) appears to be made by blending two incompatible structures, shown as (a) and (b).

i) I'm not going to solely blame all of man's activities on changes in climate
   a. I'm not going to solely blame all of climate change on man's activities
   b. I'm not going to solely blame all changes in climate on man's activities
ii) When a car seat is misused improperly
   a. When a car seat is misused
   b. When a car seat is used improperly
iii) I miss being out of touch with academia
   a. I miss academia
   b. I'm out of touch with academia

One way this can occur is by proceeding after an incompatibility arises without adjusting to it (by discarding the offending element(s)). Alternatively, as Bock and Ferreira (2014: 34) suggest, these errors can arise as a result of parallel processing, where the (incompatible) structures are separately and independently produced and not properly integrated. This sort of processing would not reflect a true all-at-once solution (or at least not a successful one), since it is not the case that the entirety of the expression is solved at once; rather, local parts of it are solved independently. Another possibility is that they result from errors in compartmentalization (e.g. the conceptual material related to 'mistakenness' is redundantly compartmentalized into both [mis-] and [improperly]).

When successful, the algorithm continues until either a satisfactory expression is formed or an incompatibility arises, in which case a candidate expression must be discarded, resetting the algorithm so it can continue again. In this way, the procedure of successively abducing and tentatively combining candidate expressions builds an expression of the original thought in language. This procedure may also be able to be enacted in an all-at-once paradigm, where the operations are carried out in parallel simultaneously, rather than successively.

As a final remark, I would like to note that the actual procedure used in expressing thought in language is likely not strictly dichotomous between these two paradigms. The extremes would be that the expression is abduced a single element at a time or else a complete expression that exhausts the entire thought is abduced all at once. But the reality may often be intermediate to these cases. In an intermediate scheme, the expression would be abduced in chunks consisting of a few elements; the process would thereby be all-at-once locally and algorithmic on a larger scale. Variation in the size of these chunks would correspond to variation in the size of the "planning units" of speech production (c.f. Konopka and Brown-Schmidt 2014: 6-10) debated in the psycholinguistic literature.

Chapter 6

Conclusion

6.1 Overall summary

This dissertation described a theory of the cognition of the expression of thought in language.

In Chapter 2, we discussed a number of background issues.

Firstly, in order for there to be a problem to be accounted for by the theory, thought cannot be conducted in the medium of language. This was the issue we took up in Section 2.1. We reviewed proposals positing that thought is conducted in language, then discussed why—despite the common intuition to the contrary—it should be assumed that thought is generally independent of language. We also brought up that there are several ways that thought may be influenced by language, in spite of this independence. We then addressed the issue of the relationship between linguistic structure and conceptual structure. I stressed the need for recognizing that there are two versions of conceptual structure that need to be considered: the conceptual structure of the original thought, and the conceptual structure depicted in the linguistic expression. We also briefly discussed how non-linguistic conceptual structure can be related to linguistic conceptual structure through conversion or evocation.

Next, in Section 2.2, we discussed the related issues of what constitutes the meaning of linguistic expressions and what the format of their mental representation is. First, we explained why lexical and grammatical aspects of linguistic meaning should be treated in an integrated manner, rather than separately. Next, we clarified that what counts as a linguistic element depends on the language user's analysis, which can vary from person to person and depending on context. After this, we discussed the granularity of linguistic meaning: the specificity or schematicity of linguistic meanings

vary and depend on the language user's analysis, and the specificity of the linguistic expression is sometimes less than that of the original thought and sometimes greater than it. Then we discussed several aspects of the content of linguistic meaning. We saw that, while linguistic meaning is commonly described as propositional, propositionality is problematic as a characterization of linguistic meaning for several reasons, and where it is apt, it still does not explain the sort of meaning held by linguistic expressions. We then saw how imagery and metaphor can be depicted in language, and discussed how metaphors in linguistic expressions reflect conceptual mappings, but cautioned against assuming that this applies to every conceivable case. The last aspect of linguistic meaning we examined was structural relationships. Following this, we considered the question of what the format of the mental representation of concepts is. We discussed the thesis that perceptual representations form the basis for the mental representation of concepts; I argued that while it is possible that perceptual representations contribute to the mental representations of concepts, it appears that perceptual representations alone cannot instantiate the mental representation of most concepts. I then introduced the topic of *conceptual lexicalization*, the cognitive ability to take any bit of conceptual material and index it as a concept (or, alternatively, the process involved in doing so). To begin with, we saw that there is great variation between languages in what concepts are lexicalized. There seems to be no restriction here, and so it appears that any arbitrary concept can be lexicalized. The process of conceptual lexicalization involves attaching an index to a unified chunk of conceptual material. The index uniquely identifies the concept and enables its retrieval. Linguistic elements are indexed concepts with a linguistic symbol attached to them (however, an indexed concept does not need to be attached to an overt form like a word); conceptual lexicalization thus supports the system of language.

In Section 2.3, we discussed the architecture of language. We reviewed several conceptions of the architecture of language. There is universal agreement that the system of language is formed by a pairing of symbol and meaning, but there is no overall consensus on the details beyond that. In the view of the architecture of language that I argued for, there is no fundamental separation of the lexicon and syntax (they are both part of a single system for expressing meaning), the relationship between the

linguistic symbol and the meaning it denotes is asymmetric (the meaning is necessarily present, but the linguistic symbol is overt in some cases and non-overt in others), and a linguistic element consists of an indexed concept with a linguistic symbol (overt or non-overt) attached to it (and accessing a linguistic element can be done without accessing an overt form, as seen in the tip-of-the-tongue phenomenon).

Finally, in Section 2.4, we outlined a schema of the process of expressing thought in language. It starts with an original language-independent thought and ends with a mental representation of a linguistic expression; our task is to describe how this expression is formulated. A linguistic expression is a 'suitable' approximation of the original thought (not necessarily the best possible approximation, since its formation may be biased by mental accessibility of linguistic structures); there is also presumably variation between individuals and situations in how closely the expression must fit the original thought in order to be acceptable. The thought that is to be expressed does not necessarily correspond to a sentence (e.g. the language user may develop the thought while formulating the linguistic expression). Also, adjustments of what to say motivated by pragmatic considerations count as reformulations of the thought, rather than reformulations of the linguistic expression (so they are outside the scope of the problem I am addressing, since such an adjustment would count as a new thought, rather than part of the expression of the thought). Since I do not assume any segregation between syntax and semantics, I do not treat syntactic encoding and semantic encoding as separate stages. Since the conceptual structure depicted by the linguistic expression is not identical to that of the original thought, expressing thought in language involves re-conceptualization. This should not be taken to mean, however, that conceptual structure is absent until it is bestowed by the formulation of the linguistic expression, or that re-conceptualization alone drives the expression of thought. Expressing thought in language is achieved by constructing an arrangement of elements that yields an approximation of the original thought. This process is essentially deriving a product representation (the linguistic expression) from a starting representation (the original thought). In order to do this, the conceptual material in the thought must be organized into available linguistic elements, and the arrangement of those elements must yield, by the mechanics of meaning construction, an approximation of the original thought.

In Chapter 3, we discussed meaning construction. The main import of meaning construction is that its mechanics underlie the fact that the meaning yielded by a particular arrangement is predictable (without requiring memorization of all of the infinite possible arrangements as individuals). First (in Section 3.1), we considered cases where this predictability may appear to break down (mismatch between the formal linguistic structure and what are actually functioning as linguistic elements and apparent underdetermination of meaning), and saw that they are not actually impediments to predictability. Then (in Section 3.2), we examined the dynamicity of meaning construction, which can be seen especially clearly in the fact that particular meanings can be approximated by entirely different expressions that take different paths to approximately the same meaning. After this (in Section 3.3), we addressed the cognitive basis of meaning construction. It appears to be related to inference. My proposal is that meaning is established through mental model cognition—manipulation of a system defined by an understanding of the function and interaction of its elements (if meaning construction is to be thought of as inferential, the relevant inferences would be made in the mental model). What I mean by 'mental model' here is distinct from Johnson-Laird's conception of mental models (since mine refers to the manipulable system, rather than a constructed iconic diorama), and my proposed account is also different from simulation semantics (since it does not require a perceptual format of mental representation, and since any simulations would have to be driven by the understanding of the function and interaction of the elements that I propose). This understanding functions as a background theory on the basis of which the meanings yielded by particular arrangements are predictable.

In Chapter 4, we discussed the issue of the construction of an arrangement of linguistic elements. I divided this issue into two intertwined facets: the organization of information in the expression and the selection of elements. First (in Section 4.1), we looked at existing theories on the organization of information in expressions; these describe some important aspects of the issue. Then I introduced the concept of *compartmentalization*, the process of organizing the conceptual material in the thought to be expressed. The thought has to be compartmentalized according to meanings in available elements that can yield a meaning approximately equivalent to the original

thought. We first saw variety in how particular chunks of conceptual material can be compartmentalized, looking at the basic variational patterns of relative shifting (changing the structural position of the element corresponding to some bit of conceptual material), relative separation/compression (distributing some bit of conceptual material into several elements or compressing the conceptual material corresponding to several elements into fewer elements), indexation versus word position (encoding constructional meaning in patterns of indexation or patterns of word position), and fundamental reformulation (achieving an approximately-equal meaning using entirely different information). We then saw how compartmentalization involves interdependencies between different constituent parts: individual bits of conceptual structure often cannot be compartmentalized in isolation; setting the compartmentalization for one bit of conceptual material affects how other bits can be compartmentalized. After this (in Section 4.2), we addressed the selection of linguistic elements. First, we discussed categorization, which is an important part of how elements are selected. Categorization is not necessarily a straightforward matter, and we saw how it works in some detail in more complex cases. Next, we discussed accessing elements, which has been extensively studied in psycholinguistics. We saw that the access of elements is in part associative and can be influenced by priming and frequency, so linguistic structures can have relatively higher or lower mental accessibility. We then saw how the issue of the selection of elements is not fully accounted for by categorization, and I proposed that the process of the selection of linguistic elements should instead be understood as a matter of compartmentalization.

In Chapter 5, we got to the core of the dissertation: my claim that the expression of thought in language is accomplished through the cognitive operation of abduction.

To begin with (in Section 5.1), we introduced abduction in general as the inference of a cause from an effect. We contrasted it with deduction (the inference of an effect from a cause) and induction (pattern extraction), and saw various kinds of problems that are solved through abduction (abductive problems are of the form *What would yield this effect?*). We then distinguished the cognitive operation of abduction (which is what my claim involves) from a conception of abduction as the navigation of a certain kind of logical system (which it does not): abduction as a cognitive operation

is the production of a cause that would yield some desired effect, through a search for cognitive objects and combinations of them; it is not inherently hypothetical, explanatory, or retrospectively evaluative.

Next (in Section 5.2), we applied abduction to the problem of expressing thought in language. We saw firstly that in a basic sense, the expression of thought in language is an abduction (the problem posed is *What would yield (an approximation of) the original thought?*); then, having noted that there are different kinds of abduction, we identified the abduction for the expression of thought as one in which the cause and the effect are approximately equal, there is no predetermined list from which the solution is taken, and there are no pre-established rules specifying what a potential solution would yield. We then discussed several aspects of this abductive cognition. Firstly, in the absence of such pre-established rules, what would be yielded by a potential solution must be predictable somehow. I proposed that mental model cognition provides this predictability by applying the understanding of how the elements of the system work as a background theory to potential linguistic elements and arrangements of them. I then pointed out that abductive problems can be solved either all at once or through multiple sequential steps, as indicated by research on insight (the psychological effect of suddenly realizing the solution to a problem). After this, we considered the issue of the search mechanism. I speculatively proposed that the search mechanism is partly based on applying the mental model to mental representations of potential elements as a sort of filter; the potential elements are automatically plugged in and activated if they go toward yielding the desired meaning. Next, we looked at the issue of how the process of expressing thought is automatized. I gave three speculative suggestions: that determining what meaning would be yielded by potential elements could be dealt with by the mechanism for accessing concepts from memory, so the yielded meaning would not have to be separately computed, that associative access can retrieve commonly-co-occurring elements, so all of the individual potential elements would not have to be looked up separately, and that prefabricated schemas of compartmentalization can be used, so the compartmentalization would not have to be figured out from scratch. Lastly, I conjectured that unconscious processing may be related to getting around working memory capacity limitations, which would otherwise present an obstacle.

Finally (in Section 5.3), we saw how the overall process of expressing thought in language could work. In the model of the process that I described, part of the thought is compartmentalized, resulting in a *candidate expression* being tentatively set; subsequent bits of the expression are abduced on the assumption of the previously-produced candidate expressions as part of the overall arrangement and incorporated into the tentative arrangement. The abductive problem is updated with the tentative addition of each successive candidate expression: *What in combination with the tentatively-set candidate expressions would yield an approximation of the thought?*. This procedure continues until either the linguistic expression is formed or an incompatibility between candidate expressions is encountered, in which case the tentatively-set arrangement is adjusted by discarding one or more candidate expressions. I proposed that this can be carried out in two paradigms: the algorithmic paradigm, which uses multiple successive steps to create the expression, and the all-at-once paradigm, which does it in a single step through parallel processing. The actual formulation of linguistic expressions may be intermediate to these two paradigms, involving the abduction of the expression in chunks of a few elements.

6.2 Remarks on some tangentially-related topics

Some ideas discussed in the dissertation are of significance to other important issues that are outside the focus of this dissertation. I will discuss (somewhat speculatively) two of these issues here.

6.2.1 Conceptual lexicalization in thought

The first topic that I wish to remark on is the role of conceptual lexicalization in thought in general. In Section 2.2.3, I discussed how conceptual lexicalization is important for language, by furnishing retrievable concepts that can be given linguistic symbols. I want to make two main points here: that the mechanism of conceptual lexicalization is used outside the domain of language, in general thinking activity, and

that it may be connected to the usage of mental representations of different levels of detail.

One of the main features of conceptual lexicalization is that indexed concepts can be referred to without any linguistic label. The idea here is that conceptual lexicalization is necessary for referring to different ideas and bits of thought, which generally do not have overt labels.

One vital function that conceptual lexicalization is involved in is the establishment of referable cognitive objects. For example, a sports coach can come up with various drills, consisting of various components, that do not have names. They can refer in their mind to this drill or that drill, even though there is no name for them. In thought, it is frequently the case that certain patterns need to be retrieved and referred to and applied in various contexts. One clear example is in computer programming (I am using programming to illustrate the principle, but it is not at all confined to the task of programming). Constructing a computer program involves using fixed operations in various combinations, and the overall problem to be solved must be broken down into sub-problems that can be solved using these operations. Often, a particular pattern will need to be applied repeatedly: it performs a particular function that is used to solve a sub-problem; it is thus a sort of 'mini-algorithm'. Whereas the basic operations are pre-existing, figuring out one of these mini-algorithms entails inventing a new thing (of course, after it is figured out, it can be applied later, which is my point). Here, the pattern of operations and the particular effect it achieves become treated as (indexed) concepts that can be referred to repeatedly: they are conceptually lexicalized. Johnson-Laird and Khemlani (2014: 27-33) give the example of figuring out how to change the position of cars on a y-shaped train track. They describe an experiment where they gave non-programmer participants various position-rearrangement tasks, and found that the participants developed algorithms of this sort to perform particular functions For example, an algorithm for reversing a train of four cars is ABCD[ ], A[BCD], [BCD]A, B[CD]A, [CD]BA, C[D]BA, [D]CBA, D[ ]CBA, [ ]DCBA, where the brackets indicate what is on the side track (empty in the initial state). These algorithms function as a sort of 'vocabulary' for solving higher-level problems (without, of course, having any explicit name).

Another important function of indexing is to enable the coordination of different ideas and trains of thought. This is used in picking up a thought that one has left off before, or connecting something one is thinking at the present moment to something one had been thinking about before; and in order to do this, of course, one needs to access the thing they were thinking about before. The same indexing is presumably necessary here. It is also used in navigating between different simultaneous threads of thought. The index is attached to the thread of thought or the idea, and uniquely identifies it.

With regard to the second point, conceptual lexicalization may be related to the construction of mental representations of different levels of detail.

Jackendoff (2007: 123-133) gives the example of making coffee. There are many different levels of detail at which the process can be conceptualized: the basic steps of [put in coffee; put in water] > [turn on coffeemaker] > [wait] can be decomposed into sub-parts at a more detailed level. For example, [put coffee in machine] consists of [prepare filter] > [[get out coffee] > [measure coffee] > [put away coffee]] > [close up filter]; [get out coffee] can be decomposed into [open freezer] > [take out coffee can] > [shut freezer], and so on. Similarly, the sports drill can be conceived of in full detail—the entire sequence and how the different parts relate to each other, which can be quite complex—or only in terms of higher functional levels, such as that it is a drill to develop some particular skill. One can have in mind 'that drill' and not have all the details of what is going on inside the drill to mentally access it as a concept. Perceptual concepts likewise can be represented at different levels of detail: the perceptual imagery can be rich or sparse. At a high level of detail, the representation consists of a rich simulation of the perceptual imagery; a low-detail representation identifies the concept and its content without actually going through the simulation. I can, for instance, refer mentally to Dvo ák's 1st symphony without running through a 'recording' of it in my head.

So there are two basic kinds of mental representation at play here. One is the *detailed representation*, in which the content is displayed in some amount of detail, and the other is the *condensed representation*, where the content is identified without being fully displayed. One example where condensed representations are explicitly 'visible' is

when a long and complex mathematical expression is defined as a single symbol; manipulation of this content can thereafter be done with reference to this symbol without referencing all of the detail of its content.

Conceptual lexicalization enables the whole of the content to be manipulated as a unit, which is necessary for condensed representations. To represent the content in toto is much more demanding on working memory than representing it in condensed form, which has a similar effect to chunking items in working memory (e.g. grouping digits of a phone number like [xxx] [xxx] [xxx]) (c.f. Levelt et al. 1999: 8; Chekaf et al. 2018). We will return to this theme in the next subsection, on the language faculty.

One issue that arises with regard to condensed and detailed representations is how the indexing mechanism works for linking the condensed representation of a concept to the content. I do not know how the indexing mechanism works, but a number of speculative possibilities are conceivable. One is that a partial 'image' of the concept can be used to then access the whole. Another is that the brain embeds a marker into mental representations that serves as something like a time-stamp indicating how long ago it was thought, imparting an intuitive sense of temporal distance. This distance may not be strictly temporal, but instead a quasi-temporal relative identification: for example, how many thoughts have occurred in the interim might also influence what the perceived temporal distance is. This mechanism can also be used to associate older bits of thought with newer bits of thought. Another possibility is that condensed representations are 'summaries' of detailed representations. Blouw et al. (2016) outline how (hypothetical) compressed mental representations called *semantic pointers* (so called because they are mental representations that point to other ones) can be used to manipulate conceptual content; these may be responsible for instantiating condensed representations.

The last speculation I will mention with regard to this topic here is that linguistic meanings may be restricted to condensed representations. Essentially any conceptual material can be indexed as a concept, and any indexed concept can be turned into a word. Now, there is linguistic thought and non-linguistic thought. Linguistic thought is of the sort that can be expressed in language; there are other kinds of thought that cannot be. For example, if a composer has a musical thought, they cannot express it in language; it has to be expressed in music. The idea of conceptual lexicalization applies

to non-linguistic thought as well. For example, emotions that correspond to particular musical structures can be indexed as a concept and referred to; so can particular patterns of notes or particular chords, etc. These concepts, like any other, can be represented in detail or in condensed form. It is true that any indexed concept can be given a linguistic label; this is a relatively trivial process. So any of these indexed musical concepts can in principle be given linguistic names (and some of them are, as terms in music theory like *V–I cadence*, etc.—although these are all very schematic concepts, and the meaning of a musical structure in a particular context in a piece is fully specific). The contents of the detailed representations of the indexed concepts are specifically musical, and therefore specifically non-linguistic. But the content that is expressed through these (would-be) linguistic names cannot be musical: it is specifically linguistic. Their representation cannot possibly be the detailed form; it has to be the condensed form. This suggests that linguistic concepts are restricted to condensed representations. Evans (2009, 2016) has made a related proposal; he distinguishes between "rich" representations (which he identifies with simulations) and "schematic" representations, and maintains that linguistic meaning is schematic and "affords access" to rich representations.

6.2.2 The language faculty

The term *language faculty* refers to the cognitive abilities that enable the use of language. What comprises the language faculty has been the topic of a considerable amount of debate, most notably in a series of papers by Hauser, Chomsky, and Fitch (Hauser et al. 2002; Fitch et al. 2005) and Pinker and Jackendoff (Pinker and Jackendoff 2005; Jackendoff and Pinker 2005). In this section, I will first give a very basic summary of the debate. I will ignore issues relating to auditory processing for speech perception, adaptations that enable vocal articulation of spoken language, and phonological structure, because these are relatively peripheral in terms of the language faculty and also not relevant to what I have to say. I will also not discuss questions like what abilities should be taken to be unique to language and what abilities should be taken to be involved in language but also used in other domains and where it is suitable to place the boundary between abilities that are part of the language faculty and those

that are not, which has occupied a substantial portion of the debate. After this, I will then add my own proposal for what constitutes the language faculty.

In the view of Chomsky and colleagues (e.g. Hauser et al. 2002; Fitch et al. 2005; Chomsky 2007; Berwick and Chomsky 2011), the core of the language faculty is the ability to take two objects and bind them into a set (called Merge) (Hauser et al. 2002 and Fitch et al. 2005 do not specifically mention Merge; they refer instead to "syntactic recursion"). This operation acts on the "conceptual atoms of the lexicon" (Berwick and Chomsky 2011: 30) and generates both syntactic embedding and syntactic displacement. This system for computation of syntactic structure interfaces with the conceptual and phonological systems, connecting syntactic objects with semantic interpretations and enabling them to be articulated. (C.f. Sections 2.2.4.1 and 2.3). They conjecture that the ability to perform Merge is the result of a single evolutionary mutation. Both Chomsky and colleagues and Pinker and Jackendoff agree that human conceptual structure is distinctive and has some place in the language faculty.

Pinker and Jackendoff (2005) argue that Merge can account for only some syntactic structure: they identify that basic devices of syntactic organization, in addition to hierarchical grouping and embedding (which can be achieved by Merge), include linear ordering (which Chomsky and colleagues take to be an artifact of externalization; see, for example Berwick and Chomsky 2011: 29), agreement, and case marking. An obvious deficiency in postulating Merge as an operation that acts on exactly two elements (also mentioned by Jackendoff 2011: 593-594) is that it would force hierarchical groupings in cases where there are none. For example, the expression *I bought apples, mangos, and pasta* would have to be given either the grouping [[[apples] [mangos]] pasta] or [apples [[mangos] [pasta]]], even if [apples], [mangos], and [pasta] are actually parallel to each other. Jackendoff (2011: 602) also points out that Merge cannot generate Semitic morphology. In infixational Semitic morphology, a root is formed from a three-consonant pattern with unspecified vowels, and is combined with patterns of vowels (with unspecified consonants) to derive specific words. Jackendoff's example is the Hebrew root *k_t_v* combined with *_a_a* (plus the affixational *–ti*) to produce *katavti* ('I wrote'). Different patterns of vowels derive different words: for example, in Egyptian Arabic, *k_t_b* can produce words like *kataba* ('he wrote'), *ka:tib*

'writing (person)', *kita:b* ('book'), and *ma-ka:tib* ('place for writing') (Aikhenvald 2007: 39); in Hebrew, *g_d_r* can produce *gadar* ('he enclosed'), *gudar* ('he was enclosed / fenced in'), *g'dor* ('enclose it!'), etc. (Bickel and Nichols 2007: 182, citing Glinert 1989). Similarly, signed languages contain morphology produced by altering the direction, speed, etc. of root signs, rather than by forming a set consisting of the root and appended alterations (Jackendoff 2011: 602). Jackendoff proposes *unification* instead as the basic combinatorial mechanism. The operation of unification achieves combination by insertion of features into slots; this would obviously account for infixational Semitic morphology and simultaneous signed morphology, and Jackendoff also suggests that it would account for the oddity of *John drank the apple* (which can only be interpreted as meaning that the apple somehow became liquid, and then John drank it): the object argument of *drink* must have the semantic feature of being liquid (ibid.: 601-602).

The suggestion I would like to put forward is that mental model cognition and conceptual lexicalization are crucially important parts of the language faculty.

As discussed in Chapter 3 (also see Chapter 5), mental model cognition as I see it is responsible for manipulating cognitive objects flexibly and combining them and determining meaning through the manipulation (with regard to mental model cognition in general, this is 'meaning' in a very broad and abstract sense; with regard to language, this is the normal sense of 'meaning'). So it enables creating meaning through arrangements of elements. It also underlies the generativity of language, where a fixed set of elements is used open-endedly to express an infinity of meanings. Mental model cognition also handles syntactic organization; computation of syntactic structure is subsumed as a special case of general arrangement of elements in a mental model.

Conceptual lexicalization, as mentioned in Section 2.2.3.2, is a necessary prerequisite for the system of denoting reusable concepts with symbols that function as 'names' for those concepts that is language: it is necessary for establishing concepts that can be retrieved repeatedly to which linguistic symbols can be attached.

The other main import of conceptual lexicalization for the language faculty is that it appears to enable condensed representations; the idea is that the ability to have condensed representations is also required for using language. Working with condensed

representations is a tremendous cognitive advantage because it allows one to manipulate hugely complex structures as single objects. The ability to manipulate information in the brain in this manner is useful outside the domain of language, of course, but evolving this sort of system may be one of the factors that enabled the advent of language. Working with multiple cognitive objects—as would be necessary, for example, in conversation—is much more easily done through condensed representations, rather than though detailed representations.

Chomsky and colleagues have made a not entirely dissimilar conjecture regarding Merge. As Boeckx (2011) articulates, the idea is that "[Merge] significantly altered Man's conceptual structures—how humans think the world. By merging lexicalized concepts, Man was able to hold in mind concepts of concepts, representations of representations, and associations of associations. *Homo* became *Homo combinans*." (p. 60). The idea I am describing here is distinct from this: their idea is that the cognitive power lies in the syntactic combination of lexicalized concepts; mine is that it lies in the concepts themselves.

6.3 Future directions

The concluding remarks of this dissertation will be on directions for future research.

One is investigating how mental model cognition works. This is a complex problem, involving determining what sort of information is stored in the brain to form what functions as the background theory of the mental model and how this information is used to implement the theory flexibly.

Another is studying how abduction works in more depth. Major issues include how creative combinations are generated in the brain, how the search mechanism functions, how certain kinds of abduction are automatized (and others not), and the role of working memory or other memory systems.

A third avenue is exploring how detailed and condensed mental representations work: what content is represented, how detailed and condensed representations of a

particular concept are connected to each other, how detailed and condensed mental representations are manipulated, and how they are neurally implemented.

Finally, the theory I laid out here can be expanded to other cognitive domains. Language is a system for expressing meaning; other meaning-bearing systems, such as mathematics, art, and music, probably share much of the mechanics of use. Some aspects will be similar, and some will not. Music, for instance, has some parallels with language (e.g. both express meaning through arrangements of elements) and some non-parallels (e.g. unlike in language, in music, the elements themselves are not actually associated with particular meanings; meaning is emergent at higher levels of organization). The cognition of musical composition involves issues like how emotional representations can be manipulated and how they relate to abstract structure. Developing and expressing ideas in mathematics and art is one of the most profound of human activities, and studying the cognition used to do it can provide an opening into understanding human creation.

References

Adams, F. and Campbell, K. (1999). Modality and abstract concepts. *Behavioral and Brain Sciences* 22: 610.

Aikhenvald, A. Y. (2007). Typological distinctions in word-formation. In T. Shopen (ed.), *Language Typology and Syntactic Description, Volume III: Grammatical Categories and the Lexicon* (Second Edition), pp. 1-65. Cambridge: Cambridge University Press.

Alario, F. X., Costa, A., and Caramazza, A. (2002). Frequency effects in noun phrase production: Implications for models of lexical access. *Language and Cognitive Processes* 17(3): 299-319.

Alderson-Day, B. and Fernyhough, C. (2015). Inner speech: development, cognitive functions, phenomenology, and neurobiology. *Psychological Bulletin* 141(5): 931-965.

Allen, A. P. and Thomas, K. E. (2011). A dual process account of creative thinking. *Creativity Research Journal* 23(2): 109-118.

Andersen, H. (1973). Abductive and deductive change. *Language* 49: 765-793.

Anderson, J. R., Anderson, J. F., Ferris, J. L., Fincham, J. M., and Jung, K. (2009). Lateral inferior prefrontal cortex and anterior cingulate cortex are engaged at different stages in the solution of insight problems. *Proceedings of the National Academy of Sciences* 106(26): 10799-10804.

Arnon, I. and Snider, N. (2010). More than words: Frequency effects for multi-word phrases. *Journal of Memory and Language* 62: 67-82.

Aronowitz, S. (2019). Memory is a modeling system. *Mind and Language* 34: 483-502.

Austin, P. (1995). Double case marking in Kanyara and Mantharta languages, Western Australia. In F. Plank (ed.), *Double case: Agreement by Suffixaufnahme*, pp. 363-79. Oxford: Oxford University Press.

Aziz-Zadeh, L., Wilson, S. M., Rizzolatti, G., and Iacoboni, M. (2006). Congruent embodied representations for visually presented actions and linguistic phrases describing actions. *Current Biology* 16: 1818-1823.

Baars, B. J. and Franklin, S. (2003). How conscious experience and working memory interact. *Trends in Cognitive Sciences* 7(4): 166-172.

Baddeley, A. (2003). Working Memory: Looking Back and Looking Forward. *Nature Reviews: Neuroscience* 4: 829-839.

Baddeley, A. (2007). *Working Memory, Thought, and Action*. Oxford: Oxford University Press.

Baddeley, A. (2012). Working Memory: Theories, Models, and Controversies. *Annual Review of Psychology* 63: 1-29.

Bamgbose, A. (1966). *A grammar of Yoruba*. Cambridge: Cambridge University Press.

Barnden, J. (1996). Unconscious Gaps in Jackendoff's 'How Language Helps Us Think'? *Pragmatics and Cognition* 4(1): 65-80.

Barnes, E. (1995). Inference to the Loveliest Explanation. *Synthese* 103(2): 251-277.

Barsalou, L. W. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences* 22: 577-660.

Barsalou, L.W. (2003). Situated simulation in the human conceptual system. *Language and Cognitive Processes* 18: 513-62.

Barsalou, L. W. (2008). Grounded cognition. *Annual Review of Psychology* 59: 617-645.

Barsalou, L.W. (2017). Cognitively Plausible Theories of Concept Composition. In J. A. Hampton and Y. Winter (eds.), *Compositionality and Concepts in Linguistics and Psychology*, pp. 9-30. Springer.

Barsalou, L. W., Santos, A., Simmons, W. K., and Wilson, C. D. (2008). Language and simulation in conceptual processing. In M. De Vega, A. M. Glenberg, and A. C. Graesser (eds.), *Symbols, embodiment, and meaning*, pp. 245–283. Oxford: Oxford University Press.

Bergen, B. K. (2012). *Louder than Words: The New Science of How the Mind Makes Meaning*. New York: Basic Books.

Bergström, F. and Eriksson, J. (2014). Maintenance of non-consciously presented information engages the prefrontal cortex. *Frontiers in Human Neuroscience* 8(938): 1-10.

Bergström, F. and Eriksson, J. (2018). Neural Evidence for Non-conscious Working Memory. *Cerebral Cortex* 28: 3217-3228.

Berlin, B., and Kay, P. (1969). *Basic color terms: Their universality and evolution*. Berkeley: University of California Press.

Berwick, R. C. (2011). Syntax Facit Saltum Redux: Biolinguistics and the Leap to Syntax. In A.M. Di Sciullo and C. Boeckx (eds.), *The Biolinguistic Enterprise: New Perspectives on the Evolution and Nature of the Human Language Faculty*, pp. 65-99. Oxford: Oxford University Press.

Berwick, R. C. and Chomsky, N. (2011). The Biolinguistic Program: The Current State of its Development. In A.M. Di Sciullo and C. Boeckx (eds.), *The Biolinguistic Enterprise: New Perspectives on the Evolution and Nature of the Human Language Faculty*, pp. 19-41. Oxford: Oxford University Press.

Bickel, B. and Nichols, J. (2007). Inflectional morphology. In T. Shopen (ed.), *Language Typology and Syntactic Description, Volume III: Grammatical Categories and the Lexicon* (Second Edition), pp. 169-240. Cambridge: Cambridge University Press.

Bickerton, D. (1995). *Language and human behavior*. Seattle, Wash.: University of Washington Press.

Bickerton, D. (2002). Language in the modular mind? It's a no-brainer! *Behavioral and Brain Sciences* 25: 21-32.

Biggam, C. P. (2012). *The Semantics of Colour: A Historical Approach*. Cambridge: Cambridge University Press.

Binder, J. R. and Desai, R. H. (2011). The Neurobiology of Semantic Memory. *Trends in Cognitive Sciences* 15(11): 527-536.

Blouw, P., Solodkin, E., Thagard, P., and Eliasmith, C. (2016). Concepts as Semantic Pointers: A Framework and Computational Model. *Cognitive Science* 40: 1128-1162.

Bock, K. and Ferreira, V. (2014). Syntactically Speaking. In M. Goldrick, V. Ferreira, and M. Miozzo (eds.), *The Oxford Handbook of Language Production*, pp. 21-46. Oxford: Oxford University Press.

Bock, K. and Griffin, Z. M. (2000). Producing words: How mind meets mouth. In L. R. Wheeldon (ed.), *Aspects of Language Production*, 7-47. Psychology Press.

Bock, K. and Loebell, H. (1990). Framing sentences. *Cognition* 35, 1-39.

Boeckx, C. (2011). Some Reflections on Darwin's Problem in the Context of Cartesian Biolinguistics. In A.M. Di Sciullo and C. Boeckx (eds.), *The Biolinguistic Enterprise: New Perspectives on the Evolution and Nature of the Human Language Faculty*, pp. 42-64. Oxford: Oxford University Press.

Boulenger, V., Hauk, O., and Pulverm ller, F. (2009). Grasping Ideas with the Motor System: Semantic Somatotopy in Idiom Comprehension. *Cerebral Cortex* 19: 1905-1914.

Bowden, E. M. and Jung-Beeman, M. (2003). Normative data for 144 compound remote associate problems. *Behavior Research Methods, Instruments, and Computers* 35(4): 634-639.

Bowden E. M., Jung-Beeman, M., Fleck, J., and Kounios, J. (2005). New approaches to demystifying insight. *Trends in Cognitive Sciences* 9(7): 322-28.

Briscoe, R. (2014). Review of *Words and Images: An Essay on the Origin of Ideas*, by Christopher Gauker. *Mind* 123(491): 902-906.

Brown, P. M. and Dell, G. S. (1987). Adapting Production to Comprehension: The Explicit Mention of Instruments. *Cognitive Psychology* 19: 441-472.

Bruening, B. (2018). The Lexicalist Hypothesis: Both Wrong and Superfluous. *Language* 94(1): 1-42.

Buechel, E. (1939). *A grammar of Lakota*. Rosebud, S. Dakota: Rosebud Educational Society.

Campos, D. (2011). On the Distinction Between Peirce's Abduction and Lipton's Inference to the Best Explanation. *Synthese* 180: 419-442.

Carey, S. (2009). *The Origin of Concepts*. Oxford: Oxford University Press.

Carey, S. (2011). Précis of *The Origin of Concepts*. *Behavioral and Brain Sciences* 34: 113-167.

Carnap, R. (1962). *Logical Foundations of Probability*. Second Edition. Chicago: University of Chicago Press.

Carruthers, P. (1996). *Language, Thought and Consciousness: An Essay In Philosophical Psychology.* Cambridge University Press.

Carruthers, P. (2002). The cognitive functions of language. *Behavioral and Brain Sciences* 25: 1-69.

Carston, R. (2002). *Thoughts and Utterances: The Pragmatics of Explicit Communication*. Oxford: Blackwell.

Casasanto, D. (2016). Temporal language and temporal thinking may not go hand in hand. In B. Lewandowska-Tomaszczyk (Ed.), *Conceptualizations of time* (pp. 169-186). Amsterdam: John Benjamins.

Casasanto, D. and Jasmin, K. (2012). The Hands of Time: Temporal gestures in English speakers. *Cognitive Linguistics* 23(4): 643-74.

Chafe, W. (1975). Some Thoughts on Schemata. In R. Schank and B. L. Nash-Webber (eds.), *Theoretical Issues in Natural Language Processing*, pp. 89-91. Cambridge, Mass.

Chafe, W. (1977). Creativity in verbalization and its implications for the nature of stored knowledge. In R. Freedle, (ed.), *Discourse Production and Comprehension*, pp. 41-55. Norwood, NJ: Ablex.

Chafe, W. (2005). The relation of grammar to thought. In C. Butler, M. d. l. A. Gomez-Gonzalez, and S. Doval-Suarez (eds.), *The Dynamics of Language Use: Functional and Contrastive Perspectives*, pp. 57-78. Amsterdam: John Benjamins.

Chatterjee, A. (2010). Disembodying cognition. *Language and Cognition* 2: 79–116.

Chekaf, M., Gauvrit, N., Guida, A., and Mathy, F. (2018). Compression in Working Memory and Its Relationship With Fluid Intelligence. *Cognitive Science* 42: 904-922.

Chomsky, N. (1957). *Syntactic Structures*. The Hague: Mouton.

Chomsky, N. (1995). *The Minimalist Program*. Cambridge, Mass.: MIT Press.

Chomsky, N. (2006). *Language and Mind*. Third Edition. Cambridge: Cambridge University Press.

Chomsky, N. (2007). Of Minds and Language. *Biolinguistics* 1: 9-27.

Clark, A. (1998). Magic Words: How Language Augments Human Computation. In P. Carruthers and J. Boucher (eds.), *Language and Thought: Interdisciplinary Themes*, pp. 162-183. Cambridge: Cambridge University Press.

Cohen L. J. (1992). *An Essay on Belief and Acceptance.* Oxford: Oxford University Press.

Comrie, B. (1981). *Language Universals and Linguistic Typology*. Chicago: University of Chicago Press.

Comrie, B. and Thompson, S. A. (2007). Lexical nominalization. In T. Shopen (ed.), *Language Typology and Syntactic Description, Volume III: Grammatical Categories and the Lexicon* (Second Edition), pp. 334-381. Cambridge: Cambridge University Press.

Corbett, G. G. (2007). Gender and noun classes. In T. Shopen (ed.), *Language Typology and Syntactic Description, Volume III: Grammatical Categories and the Lexicon* (Second Edition), pp. 241-279. Cambridge: Cambridge University Press.

Costa, A., Alario, F. X., and Caramazza, A. (2005). On the categorical nature of the semantic interference effect in the picture–word interference paradigm. *Psychonomic Bulletin & Review* 12(1): 125-131.

Cowan, N. (2008). What are the differences between long-term, short-term, and working memory? *Progress in Brain Research* 169: 323-338.

Croft, W. (2001). *Radical Construction Grammar*. Oxford: Oxford University Press.

Croft, W. (2003). *Typology and Universals*. Second Edition. Cambridge: Cambridge University Press.

Croft, W. (2007). The origins of grammar in the verbalization of experience. *Cognitive Linguistics* 18(3): 339-382.

Croft, W. (2009). Constructions and generalizations. *Cognitive Linguistics* 20(1): 157-166.

Croft, W. (2012). *Verbs: Aspect and Causal Structure*. Oxford: Oxford University Press.

Csikszentmihalyi, M. (1996). *Creativity: Flow and the Psychology of Discovery and Invention*. HarperCollins.

Cutting, J. C. and Ferreira, V. S. (1999). Semantic and Phonological Information Flow in the Production Lexicon. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 25(2): 318-344.

Dąbrowska, E. (2012). Different speakers, different grammars: Individual differences in native language attainment. *Linguistic Approaches to Bilingualism* 2: 219-253.

Dancygier, B. and Sweetser, E. (2005). *Mental Spaces in Grammar: Conditional constructions*. Cambridge: Cambridge University Press.

Damasio, A. R. (1989). Time-locked multiregional retroactivation: A systems-level proposal for the neural substrates of recall and recognition. *Cognition* 33: 25-62.

Damian, M. F. and Martin, R. C. (1999). Semantic and Phonological Codes Interact in Single Word Production. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 25(2): 334-361.

Danek, A. H., Wiley, J., and Öllinger, M. (2016). Solving Classical Insight Problems Without Aha! Experience: 9 Dot, 8 Coin, and Matchstick Arithmetic Problems. *Journal of Problem Solving* 9: 47-57.

Davies, J. (1981). As cited in Croft 2003. *Kobon*. Amsterdam: North-Holland.

Dayley, J. (1985). *Tzutujil reference grammar*. Berkeley and Los Angeles: University of California Press.

de Almeida, R. and Dwivedi, V. (2008). Coercion without Lexical Decomposition: Type-Shifting Effects Revisited. *Canadian Journal of Linguistics* 53(2): 301-326

Dehaene, S. and Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework. *Cognition* 79: 1-37.

Dehaene, S., Changeux, J.-P., Naccache, L., Sackur, J., and Sergent, C. (2006). Conscious, preconscious, and subliminal processing: a testable taxonomy. *Trends in Cognitive Sciences* 10(5): 204-211.

Dell, G. S., Schwartz, M. F., Martin, N., Saffran, E. M., and Gagnon, D. A. (1997). Lexical Access in Aphasic and Nonaphasic Speakers. *Psychological Review* 104(4): 801-838.

Dell, G. S., Nozari, N., and Oppenheim, G. M. (2014). Word Production: Behavioral and Computational Considerations. In M. Goldrick, V. Ferreira, and M. Miozzo

(eds.), *The Oxford Handbook of Language Production*, pp. 88-104. Oxford: Oxford University Press.

Della Rosa, P. A., Catricalà, E., Canini, M., Vigliocco, G., and Cappa, S. F. (2018). The left inferior frontal gyrus: A neural crossroads between abstract and concrete knowledge. *NeuroImage* 175: 449-459.

Desai, R. H., Conant, L. L., Binder, J. R., Park, H., and Seidenberg, M. (2013). A piece of the action: modulation of sensory-motor regions by idioms and metaphors. *NeuroImage* 83: 862-869.

Deutscher, G. (2002). On the misuse of the notion of 'abduction' in linguistics. *Journal of Linguistics* 38: 469-485.

Dijksterhuis, A. and Meurs, T. (2006). Where creativity resides: The generative power of unconscious thought. *Consciousness and Cognition* 15: 135-146.

Dijksterhuis, A. and Nordgren, L. F. (2006). A Theory of Unconscious Thought. *Perspectives on Psychological Science* 1(2): 95-109.

Dove, G. (2016). Three symbol ungrounding problems: Abstract concepts and the future of embodied cognition. *Psychonomic Bulletin & Review* 23: 1109-1121.

Dryer, M. S. (2007). Noun phrase structure. In T. Shopen (ed.), *Language Typology and Syntactic Description, Volume II: Complex Constructions* (Second Edition), pp. 151-205. Cambridge: Cambridge University Press.

Ellamil, M., Dobson, C., Beeman, M., and Christoff, K. (2012). Evaluative and generative modes of thought during the creative process. *NeuroImage* 59: 1783-1794.

Ellis, C. D. (1983). *Spoken Cree*. Revised Edition. Edmonton, Canada: Pica Pica Press.

Eriksson, J., Vogel, E. K., Lansner, A., Bergström, F., and Nyberg, L. (2015). Neurocognitive Architecture of Working Memory. *Neuron* 88: 33-46.

Evans, N. and Levinson, S. C. (2009). The myth of language universals: Language diversity and its importance for cognitive science. *Behavioral and Brain Sciences* 32: 429-492.

Evans, V. (2003). *The Structure of Time: Language, Meaning and Temporal Cognition*. Amsterdam: Benjamins.

Evans, V. (2005). The meaning of time: polysemy, the lexicon and conceptual structure. *Journal of Linguistics* 41(1): 33-75.

Evans, V. (2009). *How Words Mean: Lexical Concepts, Cognitive Models, and Meaning Construction*. Oxford: Oxford University Press.

Evans, V. (2013). Temporal frames of reference. *Cognitive Linguistics* 24(3): 393-435.

Evans, V. (2016). Design Features for Linguistically-Mediated Meaning Construction: The Relative Roles of the Linguistic and Conceptual Systems in Subserving the Ideational Function of Language. *Frontiers in Psychology* 7(156): 1-12.

Fauconnier, G. (1985). *Mental Spaces*. Cambridge, Mass.: MIT Press.

Fauconnier, G. (1997). *Mappings in Thought and Language*. Cambridge: Cambridge University Press.

Fauconnier, G. (2014). Mental spaces, language modalities, and conceptual integration. In M. Tomasello (ed.), *The New Psychology of Language: Cognitive and Functional Approaches to Language Structure (I)*, pp. 230-258. Psychology Press.

Fauconnier, G. and Turner, M. (1996). Blending as a Central Process of Grammar. In Adele Goldberg (ed.), *Conceptual Structure, Discourse, and Language*, pp. 113-130. Cambridge: Cambridge University Press.

Fauconnier, G. and Turner, M. (2002). *The Way We Think: Conceptual Blending and the Mind's Hidden Complexities.* New York: Basic Books.

Ferreira, V. S. (1996). Is It Better to Give Than to Donate? Syntactic Flexibility in Language Production. *Journal of Memory and Language* 35: 724-755.

Fitch, W. T., Hauser, M. D., and Chomsky, N. (2005). The evolution of the language faculty: clarifications and implications. *Cognition* 97(2): 179-210.

Flach, P. A. and Kakas, A. C. (2000). Abductive and Inductive Reasoning: Background and Issues. In P. A. Flach and A. C. Kakas (eds.), *Abduction and Induction*, pp. 1-30. Dordrecht: Springer.

Fleck, J. I. and Weisberg, R. W. (2013). Insight versus analysis: Evidence for diverse methods in problem solving. *Journal of Cognitive Psychology* 25(4): 436-463.

Fodor, J. (1975). *The Language of Thought.* New York: Thomas Y. Crowell Company.

Fodor, J. (1983). *The Modularity of Mind.* Cambridge, Mass.: MIT Press.

Fodor, J. (2001). Language, Thought and Compositionality. *Mind and Language* 16(1): 1-15.

Frankish, K. (1998). Natural Language and Virtual Belief. In P. Carruthers and J. Boucher (eds.), *Language and Thought: Interdisciplinary Themes*, pp. 248-269. Cambridge: Cambridge University Press.

Frankish, K. (2002). Language, consciousness, and cross-modular thought. *Behavioral and Brain Sciences* 25: 29-30.

Frankish, K. (2018). Inner Speech and Outer Thought. In P. Langland-Hassan and A. Vicente (eds.), *Inner Speech: New Voices*, pp. 221-243. Oxford: Oxford University Press.

Frankish, K. and Evans, J. (2009). The duality of mind: An historical perspective. In J. Evans and K. Frankish (eds.), *In Two Minds: Dual Processes and Beyond*, pp. 1-29. Oxford: Oxford University Press.

Friederici, A. D. (2002). Towards a neural basis of auditory sentence processing. *Trends in Cognitive Sciences* 6(2): 78-84.

Friederici, A. D., Mueller, J. L., and Oberecker, R. (2011). Precursors to Natural Grammar Learning: Preliminary Evidence from 4-Month-Old Infants. *PLoS ONE* 6(3): 1-7.

Gaby, A. and Sweetser, E. (2017). Space-time Mappings Beyond Language. In B. Dancygier (ed.), *The Cambridge Handbook of Cognitive Linguistics*, pp. 635-650. Cambridge: Cambridge University Press.

Galetzka, C. (2017). The Story So Far: How Embodied Cognition Advances Our Understanding of Meaning-Making. *Frontiers in Psychology* 8(1315): 1-5.

Gallese, V. and Lakoff, G. (2005). The brain's concepts: the role of the sensory-motor system in conceptual knowledge. *Cognitive Neuropsychology* 22; 455-479.

Gauker, C. (1994). *Thinking Out Loud: An Essay on the Relation Between Thought and Language*. Princeton, New Jersey: Princeton University Press.

Gauker, C. (2002). No conceptual thought without language. *Behavioral and Brain Sciences* 25: 31.

Gauker, C. (2011). *Words and Images: An Essay on the Origin of Ideas*. Oxford and New York: Oxford University Press.

Gauker, C. (2018). Inner Speech as the Internalization of Outer Speech. In P. Langland-Hassan and A. Vicente (eds.), *Inner Speech: New Voices*, pp. 53-77. Oxford: Oxford University Press.

Ghiselin, B. (1952). *The Creative Process: A Symposium*. Berkeley and Los Angeles: University of California Press.

Gibbs, R. W. (2009). Why Do Some People Dislike Conceptual Metaphor Theory? *Cognitive Semiotics* 5(1-2): 14-36.

Gibbs, R. W. (2011). Evaluating Conceptual Metaphor Theory. *Discourse Processes* 48: 529-562.

Gibbs, R. W. and Colston, H. L. (1995). The cognitive psychological reality of image schemas and their transformations. *Cognitive Linguistics* 6(4): 347-378.

Gilhooly, K. J. (2016). Incubation and Intuition in Creative Problem Solving. *Frontiers in Psychology* 7(1076): 1-9.

Givón, T. (1979). *On understanding grammar*. New York: Academic Press.

Glenberg, A. M. and Kaschak, M. P. (2002). Grounding language in action. *Psychonomic Bulletin & Review* 9(3): 558-565.

Glenberg, A. M. and Gallese, V. (2012). Action-based language: A theory of language acquisition, comprehension, and production. *Cortex* 48: 905-922.

Glinert, L. (1989). *The grammar of Modern Hebrew*. Cambridge: Cambridge University Press.

Goldberg, A. E. (1995). *Constructions: A construction grammar approach to argument structure*. University of Chicago Press.

Goldberg, A. E. (2006). *Constructions at work: The nature of generalization in language*. Oxford: Oxford University Press.

Goldberg, A. E. (2009). The nature of generalization in language. *Cognitive Linguistics* 20(1): 93-127.

Goldrick, M. (2014). Phonological Processing: The Retrieval and Encoding of Word Form Information in Speech Production. In M. Goldrick, V. Ferreira, and M. Miozzo (eds.), *The Oxford Handbook of Language Production*, pp. 228-244. Oxford: Oxford University Press.

González, J., Barros-Loscertales, A., Pulvermüller, F., Meseguer, V., Sanjuán, A., Belloch, V., and Avila, C. (2006). Reading cinnamon activates olfactory brain regions. *NeuroImage* 32(2): 906-912.

Grady, J. E. (1998). The "conduit metaphor" revisited: A reassessment of metaphors for communication. *Discourse and cognition: Bridging the gap*, pp. 205-218.

Grady, J. E. (2005). Image schemas and perception: refining a definition. In B. Hampe (ed.), *From perception to meaning: image schemas in cognitive linguistics*, pp. 35-56. Berlin: Mouton de Gruyter.

Grady, J, Oakley, T., and Coulson, S. (1999). Conceptual Blending and Metaphor. In G. Steen and R. Gibbs (eds.), *Metaphor in cognitive linguistics*, pp. 101-124. Amsterdam and Philadelphia: John Benjamins.

Grady, J., Taub, S., and Morgan, P. (1996). Primitive and compound metaphors. In A. Goldberg (ed.), *Conceptual Structure, Discourse and Language*, pp. 177-187. Stanford, Calif: CSLI Publications.

Green, G. (1973) A Syntactic Syncretism in English and French. In B. Kachru, R. B. Lees, Y. Malkiel, A. Pietrangeli, and S. Saporta (eds.), *Issues in Linguistics*, pp. 257-278. Urbana, Ill.: University of Illinois Press.

Gumperz, J. J. (1982). *Discourse strategies*. Cambridge: Cambridge University Press.

Hadamard, J. (1945). *The Psychology of Invention in the Mathematical Field*. New York: Princeton University Press and Dover Publications.

Haiman, J. (1980). Dictionaries and encyclopedias. *Lingua* 50: 329-57.

Haimovici, S. (2018). The Modal–Amodal Distinction in the Debate on Conceptual Format. *Philosophies* 3(7): 1-13. doi:10.3390/philosophies3020007

Hampton, J. A. (2017). Compositionality and Concepts. In J. A. Hampton and Y. Winter (eds.), *Compositionality and Concepts in Linguistics and Psychology*, pp. 95-122. Springer.

Hampton, J. A. and Winter, Y. (eds.) (2017). *Compositionality and Concepts in Linguistics and Psychology*. Springer.

Hardin, C. L. and Maffi, L. (eds.) (1997). *Color Categories in Thought and Language*. Cambridge: Cambridge University Press.

Harman, G. H. (1965). The inference to the best explanation. *Philosophical Review* 74: 88-95.

Harnad, S. (1990). The Symbol Grounding Problem. *Physica D* 42: 335-346.

Haspelmath, M. (2007). Coordination. In T. Shopen (ed.), *Language Typology and Syntactic Description, Volume II: Complex Constructions* (Second Edition), pp. 1-51. Cambridge: Cambridge University Press.

Hartsuiker, R. J. (2014). Monitoring and Control of the Production System. In M. Goldrick, V. Ferreira, and M. Miozzo (Eds.), *The Oxford Handbook of Language Production*, pp. 417-346. Oxford: Oxford University Press.

Hartsuiker, R. J., Pickering, M. J., and Veltkamp, E. (2004). Is Syntax Separate or Shared Between Languages? Cross-Linguistic Syntactic Priming in Spanish-English Bilinguals. *Psychological Science* 15(6): 409-414.

Haser, V. (2005). *Metaphor, metonymy, and experientialist philosophy: Challenging cognitive semantics*. Berlin: Mouton de Gruyter.

Hauk, O., Johnsrude, I., and Pulvermüller, F. (2004). Somatotopic representation of action words in human motor and premotor cortex. *Neuron* 41: 301-307.

Hauk, O. and Tschentscher, N. (2013). The body of evidence: what can neuroscience tell us about embodied semantics? *Frontiers in Psychology* 4(50): 1-14.

Hauser, M. D., Chomsky, N., and Fitch, W. T. (2002). The faculty of language: what is it, who has it, and how did it evolve? *Science* 298: 1569-1579.

Hassin, R. R., Bargh, J. A., Engell, A. D., and McChulloch, K. C. (2009). Implicit working memory. *Consciousness and Cognition* 18: 665-678.

Haviland, S. and Clark, H. (1974). What's New? Acquiring New Information as a Process in Comprehension. *Journal of Verbal Learning and Verbal Behavior* 13: 512-521.

Hayase, N. (2018). Reconsidering the Status of Constructional Schema. *Journal of Cognitive Linguistics: The Journal of the Japanese Cognitive Linguistics Association* 3: 72-81.

Heine, B., Claudi, U., and H nnemeyer, F. (1991). *Grammaticalization: A Conceptual Framework*. Chicago: University of Chicago Press.

Hinzen, W. (2006). *Mind Design and Minimal Syntax*. Oxford: Oxford University Press.

Hinzen, W. (2011). Emergence of a Systemic Semantics through Minimal and Underspecified Codes. In A.M. Di Sciullo and C. Boeckx (eds.), *The Biolinguistic Enterprise: New Perspectives on the Evolution and Nature of the Human Language Faculty*, pp. 417-439. Oxford: Oxford University Press.

Holton, D., Mackridge, P., and Philippaki-Warburton, I. (1997). *Greek: a comprehensive grammar of the modern language*. London: Routledge.

Hopper, P. J. and Traugott, E. (2003). *Grammaticalization*. Cambridge: Cambridge University Press.

Hurlburt, R. T. and Akhter, S. A. (2008). Unsymbolized thinking. *Consciousness and Cognition* 17: 1364-1374.

Hurlburt, R.T., Heavey, C. L., and Kelsey, J. M. (2013). Toward a phenomenology of inner speaking. *Consciousness and Cognition* 22: 1477-1494.

Indefrey, P. and Levelt, W. (2004). The spatial and temporal signatures of word production components. *Cognition* 92: 101-144.

Jackendoff, R. (1996). How language helps us think. *Pragmatics and Cognition* 4(1): 1-34.

Jackendoff, R. (1997). *The Architecture of the Language Faculty*. Cambridge, Mass.: MIT Press.

Jackendoff, R. (2002). *Foundations of Language*. Oxford and New York: Oxford University Press.

Jackendoff, R. (2007). *Language, Consciousness, Culture: Essays on Mental Structure*. Cambridge, Mass.: MIT Press.

Jackendoff, R. (2011). What is the human language faculty? Two views. *Language* 87(3): 586-624.

Jackendoff, R. (2012). *A User's Guide to Thought and Meaning*. Oxford and New York: Oxford University Press.

Jackendoff, R. (2017). In Defense of Theory. *Cognitive Science* 41: 185-212.

Jackendoff, R. and Pinker, S. (2005). The nature of the language faculty and its implications for evolution of language (Reply to Fitch, Hauser, and Chomsky). *Cognition* 97: 211-225.

Jacobs, C. L., Cho, S., and Watson, D. G. (2019). Self-Priming in Production: Evidence for a Hybrid Model of Syntactic Priming. *Cognitive Science* 43: e12749.

Janssen, N. and Barber, H. A. (2012). Phrase Frequency Effects in Language Production. *PLoS ONE* 7(3): 1-11.

Johnson, M. (1987). *The Body in the Mind: The Bodily Basis of Meaning, Imagination and Reason*. Chicago: University of Chicago Press.

Johnson, M. (2007). The meaning of the body: Aesthetics of human understanding. Chicago: University of Chicago Press.

Johnson-Laird, P. N. (1980). Mental Models in Cognitive Science. *Cognitive Science* 4: 71-115.

Johnson-Laird, P. N. (1983). Mental models. Cambridge: Cambridge University Press.

Johnson-Laird, P. N. (2010). Mental models and human reasoning. *Proceedings of the National Academy of Sciences* 107(43): 18243-18250.

Johnson-Laird, P. N. and Byrne, R. M. J. (2002). Conditionals: A Theory of Meaning, Pragmatics, and Inference. *Psychological Review* 109(4): 646-678.

Johnson-Laird, P. N. and Khemlani, S. S. (2014). Toward a Unified Theory of Reasoning. *Psychology of Learning and Motivation* 59: 1-42.

Jönsson, M. L. (2015). Overextension in verb conjunctions. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 41(6): 1917-1922.

Jung-Beeman, M., Bowden, E. M., Haberman, J., Frymiare, J. L., Arambel-Liu, S., Greenblatt, R., Reber, P. J., and Kounios, J. (2004). Neural Activity When People Solve Verbal Problems with Insight. *PLoS Biology* 2(4): 500-510.

Kahneman, D. (2011). *Thinking, Fast and Slow*. Farrar, Straus and Giroux.

Kaschak, M. P., Jones, J. L., Carranza, J., and Fox, M. R. (2014). Embodiment and Language Comprehension. In Lawrence Shapiro (ed.), *The Routledge Handbook of Embodied Cognition*, pp. 118-126. London and New York: Routledge.

Kay, P., Brent, B., Maffi, L., Merrifield, W. R., and Cook, R. (2009). *The World Color Survey*.

Keller, C. M., and Keller, J. D. (1996). Imaging in Iron, or Thought is not Inner Speech. In J. J. Gumperz and S. C. Levinson (eds.), *Rethinking linguistic relativity*, pp. 115-129. Cambridge: Cambridge University Press.

Kiefer, M. and Pulvermüller, F. (2012). Conceptual representations in mind and brain: Theoretical developments, current evidence, and future directions. *Cortex* 48: 805-825.

Kiefer, M., Sim, E.-J., Herrnberger, B., Grothe, J., and Hoenig, K. (2008). The sound of concepts: Four markers for a link between auditory and conceptual brain systems. *Journal of Neuroscience* 28: 12224-12230.

Kita, S. and Özyürek, A. (2003). What does cross-linguistic variation in semantic coordination of speech and gesture reveal? Evidence for an interface representation of spatial thinking and speaking. *Journal of Memory and Language* 48(1): 16-32.

Knoblich, G., Ohlsson, S., Haider, H., and Rhenius, D. (1999). Constraint relaxation and chunk decomposition in insight problem solving. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 25(6): 1534-1555.

Koestler, A. (1964). *The Act of Creation*. London: Hutchinson & Co.

Konopka, A. E. and Brown-Schmidt, S. (2014). Message Encoding. In M. Goldrick, V. Ferreira, and M. Miozzo (Eds.), *The Oxford Handbook of Language Production*, pp. 3-20. Oxford: Oxford University Press.

Kossmann, M. G. (2000). As cited in Dryer 2007. *Esquisse grammaticale du rifain oriental*. Paris: Peeters.

Kouider, S. and Dehaene, S. (2007). Levels of processing during non-conscious perception: a critical review of visual masking. *Philosophical Transactions of the Royal Society B* 363: 857-875.

Kounios, J. and Beeman, M. (2014). The Cognitive Neuroscience of Insight. *Annual Review of Psychology* 65: 71-93.

Kounios, J., Frymiare, J. L., Bowden, E. M., Fleck, J. I., Subramaniam, K., Parrish, T. B., and Jung-Beeman, M. (2006). The Prepared Mind: Neural Activity Prior to Problem Presentation Predicts Subsequent Solution by Sudden Insight. *Psychological Science* 17(10): 882-890.

Kövecses, Z. (2015). *Where Metaphors Come from: Reconsidering Context in Metaphor*. New York: Oxford University Press.

Kraft, C. H. and Kirk-Green, A. H. M. (1973). *Hausa*. London: English Universities Press.

Ladyman, J. (2002). *Understanding Philosophy of Science*. London and New York: Routledge.

Ladyman, J., Douven, I., Horsten, L., and van Fraassen, B. (1997). A defence of van Fraassen's critique of abductive inference: reply to Psillos. *The Philosophical Quarterly* 47(188): 305-321.

Lakoff, G. (1987). *Women, Fire, and Dangerous Things: What Categories Reveal about the Mind*. Chicago and London: The University of Chicago Press.

Lakoff, G. (1993). The contemporary theory of metaphor. In A. Ortony (ed.), *Metaphor and Thought*, pp. 202-251. Cambridge: Cambridge University Press.

Lakoff, G. (2008). The neural theory of metaphor. In R. W. Gibbs, Jr. (ed.), *The Cambridge handbook of metaphor and thought*, pp. 17-38. Cambridge: Cambridge University Press.

Lakoff, G. and Johnson, M. (1980). *Metaphors We Live By*. Chicago: The University of Chicago Press.

Lakoff, G. and Johnson, M. (1999). *Philosophy in the flesh: The embodied mind and its challenge to western thought*. New York: Basic books.

Lambon Ralph, M. A. (2014). Neurocognitive insights on conceptual knowledge and its breakdown. *Philosophical Transactions of the Royal Society B* 369: 1-11.

Lambon Ralph, M. A., Pobric, G., and Jefferies, E. (2009). Conceptual knowledge is underpinned by the temporal lobe bilaterally: Convergent evidence from rTMS. *Cerebral Cortex* 19: 832-838.

Langacker, R. W. (1987). *Foundations of Cognitive Grammar, Volume I: Theoretical Prerequisites*. Stanford, Calif.: Stanford University Press.

Langacker, R. W. (1993). Reference-Point Constructions. *Cognitive Linguistics* 4: 1-38.

Langacker, R. W. (1997). The contextual basis of cognitive semantics. In J. Nuyts and E. Pederson (eds.), *Language and conceptualization*, pp. 229-252.

Langacker, R. W. (1999). Virtual Reality. *Studies in the Linguistic Sciences* 29(2):77-103.

Langacker, R. W. (2000). *Grammar and Conceptualization*. Berlin/New York: Mouton de Gruyter.

Langacker, R. W. (2002). Discourse in Cognitive Grammar. *Cognitive Linguistics* 12(2): 143-188.

Langacker, R. W. (2003). Constructions in Cognitive Grammar. *English Linguistics* 20(1): 41-83.

Langacker, R. W. (2008). *Cognitive Grammar: A Basic Introduction*. Oxford and New York: Oxford University Press.

Langacker, R. W. (2009a). Constructions and constructional meaning. In V. Evans and S. Pourcel (eds.), *New directions in cognitive linguistics*, pp. 225-267. Amsterdam and Philadelphia: John Benjamins Publishing Company.

Langacker, R. W. (2009b). *Investigations in Cognitive Grammar*. Berlin/New York: Mouton de Gruyter.

Langacker, R. W. (2017). Entrenchment in cognitive grammar. In H.-J. Schmid (ed.), *Entrenchment and the psychology of language learning: How we reorganize and adapt linguistic knowledge*, pp. 39-56. Washington, D.C. and Berlin: American Psychological Association and De Gruyter Mouton.

Langland-Hassan, P. (2018). From Introspection to Essence: The Auditory Nature of Inner Speech. In P. Langland-Hassan and A. Vicente (eds.) *Inner Speech: New Voices*. Oxford University Press.

Leshinskaya, A. and Caramazza, A. (2014). Nonmotor Aspects of Action Concepts. *Journal of Cognitive Neuroscience* 26(12): 863-2879.

Levelt, W. (1989). *Speaking: From intention to articulation.* Cambridge, Mass.: MIT Press.

Levelt, W., Roelofs, A., and Meyer, A. (1999). A theory of lexical access in speech production. *Behavioral and Brain Sciences* 22: 1-75.

Levin, B. and Rappaport Hovav, M. (2019). Lexicalization Patterns. In R. Truswell (ed.), *The Oxford Handbook of Event Structure*, pp. 395-425. Oxford: Oxford University Press.

Levinson, S. C. (1997). From outer to inner space: linguistic categories and non-linguistic thinking. In J. Nuyts and E. Pederson (eds.), *Language and conceptualization*, pp. 13-45. Cambridge: Cambridge University Press.

Levinson, S. C. (2003). *Space in Language and Cognition: Explorations in Cognitive Diversity*. Cambridge: Cambridge University Press.

Li, C. and Thompson, S. A. (1981). *Mandarin Chinese: a functional reference grammar*. Berkeley and Los Angeles: University of California Press.

Lipton, P. (2004). *Inference to the Best Explanation*. Second Edition. New York: Routledge.

Loebell, H. and Bock, K. (2003). Structural priming across languages. *Linguistics* 41(5): 791-824.

MacDonald, L. W., Biggam, C. P., Maramei, G. V. (eds.) (2018). *Progress in Colour Studies: Cognition, Language, and Beyond.* Amsterdam and Philadelphia: John Benjamins Publishing Company.

Machery, E. (2005). You Don't Know How You Think: Introspection and Language of Thought. *British Journal for the Philosophy of Science* 56: 469-485.

Machery, E. (2007). Concept empiricism: A methodological critique. *Cognition* 104: 19-46.

Machery, E. (2009). *Doing without Concepts*. Oxford: Oxford University Press.

Machery, E. (2010). Précis of *Doing without Concepts*. *Behavioral and Brain Sciences* 33: 195-244.

Machery, E. (2012). Review of *Words and Images: An Essay on the Origin of Ideas*, by Christopher Gauker. *Notre Dame Philosophical Reviews*. https://ndpr.nd.edu/news/words-and-images-an-essay-on-the-origin-of-ideas/

Machida, A. [町田章]. (2012). Subjectivity and Making Invisible Participants Visible: The Cognitive Process of Objectification [主観性と見えない参与者の可視化―客体化の認知プロセス―]. *Papers from the National Conference of the Japanese Cognitive Linguistics Association* [日本認知言語学会論文集] 12: 246-258.

Machida, A. [町田章]. (2013). Embodied Experiencers and Observers: The Limits of the Stage Model [身体的経験者と観察者－ステージモデルの限界－]. *Papers from the National Conference of the Japanese Cognitive Linguistics Association* [日本認知言語学会論文集] 13: 661-666.

Machida, A. (2017). On "Objective" Construals: A Cognitive Grammar Account of (Inter)subjectification. Paper presented at the 14th International Cognitive Linguistics Conference, Tartu, Estonia.

Machida, A. [町田章]. (2020). A typology of intersubjectivity and grounding: focusing on so-called argument omission [間主観性の類型とグラウンディング―いわゆる項の省略現象を中心に―]. In S. Minami [南佑亮], T. Honda [本田隆裕], and E. Tanaka [田中英理] (eds.), *Deepening and Expanding from English Linguistics* [英語学の深まり・英語学からの広がり], pp. 245-258. Tokyo: Eihosha.

Madsen, M. W. (2016). Cognitive Metaphor Theory and the Metaphysics of Immediacy. *Cognitive Science* 40: 881-908.

Magee, B. (1997). *Confessions of a Philosopher: A Personal Journey through Western Philosophy, from Plato to Popper*. London: Weidenfeld and Nicolson.

Magnani, L. (2009). *Abductive Cognition: The Epistemological and Eco-Cognitive Dimensions of Hypothetical Reasoning*. Berlin: Springer.

Mahon, B. Z. (2015a). What is embodied about cognition? *Language, Cognition and Neuroscience* 30(4): 420-429. Manuscript (downloaded on June 30, 2020 from https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4405253/pdf/nihms-644236.pdf).

Mahon, B. Z. (2015b). The Burden of Embodied Cognition. *Canadian Journal of Experimental Psychology* 69(2): 172-178.

Mahon, B. Z. and Hickok, G. (2016). Arguments about the nature of concepts: Symbols, embodiment, and beyond. *Psychonomic Bulletin & Review* 23: 941-958.

Malt, B. C. and Majid, A. (2013). How thought is mapped into words. *Wiley Interdisciplinary Reviews: Cognitive Science* 4(6): 583-597.

Mandler, J. M. (2004). Thought before language. *Trends in Cognitive Sciences* 8(11): 508-513.

Mandler, J. M. and Pagán Cánovas, C. (2014). On defining image schemas. *Language and Cognition* 6(4): 1-23.

Margolis, E. and Laurence, S. (eds.) (2015). *The Conceptual Mind: New Directions in the Study of Concepts*. Cambridge, Mass.: MIT Press.

Markman, A. B., and Stillwell, C. H. (2004). Concepts á la modal: an extended review of Prinz's *Furnishing the mind*. *Philosophical Psychology* 17(3): 391-401.

McAuliffe, W. (2015). How Did Abduction Get Confused with Inference to the Best Explanation? *Transactions of the Charles S. Peirce Society* 51: 300-319.

McGlone, M. S. (2007). What is the explanatory value of a conceptual metaphor? *Language and Communication* 27: 109-126.

McGlone, M. S. (2011). Hyperbole, Homunculi, and Hindsight Bias: An Alternative Evaluation of Conceptual Metaphor Theory. *Discourse Processes* 48(8): 563-574.

McGlone, M. S. and Harding, J. L. (1998). Back (or Forward?) to the Future: The Role of Perspective in Temporal Language Comprehension. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 24(5): 1211-1223.

Metcalfe, J. (1986). Premonitions of insight predict impending error. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 12: 623-634.

Meyer, A. S. (1996). Lexical Access in Phrase and Sentence Production: Results from Picture–Word Interference Experiments. *Journal of Memory and Language* 35: 477-496.

Meyer, D. E. and Schvaneveldt, R. W. (1971). Facilitation in Recognizing Pairs of Words: Evidence of a Dependence between Retrieval Operations. *Journal of Experimental Psychology* 90(2): 227-234.

Miller, M. B., Donovan, C-L, Bennett, C. M., Aminoff, E. M., and Mayer, R. E. (2012). Individual differences in cognitive style and strategy predict similarities in the patterns of brain activity between individuals. *NeuroImage* 59: 83-93.

Minervino, R. A., Martin, A., Tavernini, L. M., and Trench, M. (2018). The Understanding of Visual Metaphors by the Congenitally Blind. *Frontiers in Psychology* 9(1242): 1-11.

Mitchell, R. W. and Clement, C. A. (1999). Simulations, simulators, amodality, and abstract terms. *Behavioral and Brain Sciences* 22: 628-629.

Mithun, M. (1984). How to avoid subordination. *Berkeley Linguistics Society* 10: 493-509.

Molina, C. (2002). Could you *think* Carruthers's ideas without having to *speak* them? *Talk* with yourself if you want to have any *thought* on that. *Behavioral and Brain Sciences* 25: 36-37.

Montefinese, M. (2019). Semantic representation of abstract and concrete words: a minireview of neural evidence. *Journal of Neurophysiology* 121: 1585-1587.

Moore, K. E. (2014). *The Spatial Language of Time*. Amsterdam and Philadelphia: John Benjamins Publishing Company.

Nonato, R. (2014). *Clause Chaining, Switch Reference and Coordination.* Ph.D. dissertation, Massachusetts Institute of Technology.

Núñez, R. and Sweetser, E. (2006). With the future behind them: Convergent evidence From Aymara language and gesture in the crosslinguistic comparison of spatial construals of time. *Cognitive Science* 30: 401-450.

Nuyts, J. (2012). Language is an instrument for thought. Really? *Pragmatics and Cognition* 20(2): 317- 333.

Ohlsson, S. (1999). Selecting is not abstracting. *Behavioral and Brain Sciences* 22: 630-631.

Oldfield, R. C. and Wingfield, A. (1965). Response latencies in naming objects. *Quarterly Journal of Experimental Psychology* 17: 273-81.

Oppenheim, G. M., Dell, G. S., and Schwartz, M. F. (2010). The dark side of incremental learning: A model of cumulative semantic interference during lexical access in speech production. *Cognition* 114: 227-252.

Ormerod, R. J. (2010). Rational inference: deductive, inductive and probabilistic thinking. *Journal of the Operational Research Society* 61: 1207-1223.

Osada, T. (1992). *A reference grammar of Mundari*. Institute for the Study of Languages and Cultures of Asia and Africa, Tokyo University of Foreign Studies

Osherson, D. N. and Smith, E. E. (1981). On the adequacy of prototype theory as a theory of concepts. *Cognition* 9: 35-58.

Öztekin, I., Davachi, L., and McElree, B. (2010). Are Representations in Working Memory Distinct From Representations in Long-Term Memory? Neural Evidence in Support of a Single Store. *Psychological Science* 21(8): 1123-1133.

Parasuraman, R., and Jiang, Y. (2012). Individual differences in cognition, affect, and performance: Behavioral, neuroimaging, and molecular genetic approaches. *NeuroImage* 59:70-82.

Partee, B. H. (2004). Compositionality and coercion in semantics: The dynamics of adjective meaning. In G. Bouma, I. Krämer, and J. Zwarts (eds.), *Cognitive Foundations of Interpretation*, pp. 145-162. Amsterdam: Royal Netherlands Academy of Arts and Sciences.

Patterson, K., Nestor, P. J., and Rogers, T. T. (2007). Where do you know what you know? The representation of semantic knowledge in the human brain. *Nature Reviews: Neuroscience* 8: 976-987.

Pecher, D. (2013). The perceptual representation of mental categories. In D. Reisberg (ed.), *The Oxford Handbook of Cognitive Psychology*, pp. 358-373. Oxford: Oxford University Press.

Pecher, D. (2018). Curb Your Embodiment. *Topics in Cognitive Science* 10: 501-517.

Pecher, D. and Zeelenberg, R. (2018). Boundaries to grounding abstract concepts. *Philosophical Transactions of the Royal Society B* 373: 1-6.

Peirce, C. S. (1931-1935). C. Hartshorne and P. Weiss (eds.), *The Collected Papers of Charles Sanders Peirce*. Cambridge, Mass.: Harvard University Press.

Pelletier, F. J. (2017). Compositionality and Concepts—A Perspective from Formal Semantics and Philosophy of Language. In J. A. Hampton and Y. Winter (eds.), *Compositionality and Concepts in Linguistics and Psychology*, pp. 31-94. Springer.

Persuh, M., LaRock, E., and Berger, J. (2018). Working Memory and Consciousness: The Current State of Play. *Frontiers in Human Neuroscience* 12(78): 1-11.

Pickering, M. J. and Garrod, S. (2013). An integrated theory of language production and comprehension. *Behavioral and Brain Sciences* 36(4): 329-347.

Pinker, S. (1994). *The Language Instinct: How the Mind Creates Language*. New York: HarperCollins and William Morrow and Company.

Pinker, S. and Jackendoff, R. (2005). The faculty of language: what's special about it? *Cognition* 95: 201-236.

Poincaré, H. (1910). Mathematical Creation. *The Monist* 20(3): 321-335.

Polo-Sherk, R. (2016a). The Process of Meaning Construction: The Parameters of *Realisity* and *Generality* [意味構築のプロセスーrealis 性と一般性のパラメターー]. In *Proceedings of the Fortieth Annual Meeting of the Kansai Linguistic Society* (KLS 36), pp. 183-194.

Polo-Sherk, R. (2016b). *The Structure of Mental Representations of Linguistic Expressions: Realisity and Generality as Parameters* [言語表現の心的表象の構造ー realis 性と一般性をパラメターとしてー]. M.A. Thesis, Hiroshima University.

Polo-Sherk, R. (2020). Compartmentalization: The Organization of Thought for Expression [言語化のための思考整理　―思考の区画化に関する提案―]. In *Papers from the 20th National Conference of the Japanese Cognitive Linguistics Association*, pp. 232-244.

Poppe, N. (1974). *Grammar of written Mongolian*. Wiesbaden: Otto Harrassowitz.

Prinz, J. (2002). *Furnishing the mind: concepts and their perceptual basis*. Cambridge, Mass.: MIT Press.

Pulvermüller, F., Hauk, O., Nikulin, V. V., and Ilmoniemi, R. J. (2005). Functional links between motor and language systems. *European Journal of Neuroscience* 21(3): 793-797.

Pylkkänen, L. and McElree, B. (2006). The syntax-semantics interface: On-line composition of sentence meaning. In M. Traxler and M.A. Gernsbacher (eds.), *Handbook of psycholinguistics* (Second Edition), pp. 537-577. New York: Elsevier.

Ran, B. and Duimering, P. R. (2010). Conceptual Combination: Models, Theories and Controversies. *International Journal of Cognitive Linguistics* 1(1): 65-90.

Raz, S. (1983). *Tigre grammar and texts.* Malibu, Calif.: Undena Press.

Reddy, M. (1979). The Conduit Metaphor—A Case of Frame Conflict in Our Language about Language. *Metaphor and thought* 2: 285-324.

Roelofs, A. (1992). A spreading-activation theory of lemma retrieval in speaking. *Cognition* 42: 107-142.

Rogers, T. T., Lambon Ralph, M. A., Garrard, P., Bozeat, S., McClelland, J. L., Hodges, J. R., and Patterson, K. (2004). Structure and deterioration of semantic memory: a neuropsychological and computational investigation. *Psychological Review* 111(1): 205-235.

Rowlands, E. C. (1969). As cited in Croft 2003. *Yoruba*. Sevenoaks: Hodder and Stoughton.

Saji, N. Imai, M., Saalbach, H., Zhang, Y., Shu, H., and Okada, H. (2011). *Cognition* 118: 45-61.

Saji, N., Saalbach, H., Imai, M., Zhang, Y., Shu, H., and Okada, H. (2008). Fast-mapping and Reorganization: Development of Verb Meanings as a System. *Proceedings of the Annual Meeting of the Cognitive Science Society* 30(30): 35-40.

Sakama, C. and Inoue, K. (2016). Abduction, conversational implicature and misleading in human dialogues. *Logic Journal of the IGPL* 24(4): 526-541.

Schmid, H.-J. (2015). A blueprint of the entrenchment-and-conventionalization model. *Yearbook of the German Cognitive Linguistics Association* 3(1): 3-26.

Schultze-Berndt, E. (2000). *Simple and complex verbs in Jaminjung: A study of event categorization in an Australian language*. Ph.D. dissertation, Radboud University. MPI Series in Psycholinguistics.

Schurz, G. (2008). Patterns of Abduction. *Synthese* 164: 201-234.

Selcan, Z. (1998). As cited in Haspelmath 2007. *Grammatik der Zaza-Sprache, Nord-Dialekt (Dersim-Dialekt)*. Berlin: Wissenschaft- und Technik-Verlag.

Shallice, T. and Cooper, R. (2013). Is there a semantic system for abstract words? *Frontiers in Human Neuroscience* 7(175): 1-10.

Shen, W. Tong, Y., Li, F., Yuan, Y., Hommel, B., Liu, C., and Luo, J. (2018). Tracking the neurodynamics of insight: A meta-analysis of neuroimaging studies. *Biological Psychology* 138: 189-198.

Sio, U. N. and Ormerod, T. C. (2009). Does Incubation Enhance Problem Solving? A Meta-Analytic Review. *Psychological Bulletin* 135(1): 94-120.

Slezak, P. (2002a). Thinking about thinking: Language, thought and introspection. *Language and Communication* 22: 353–373.

Slezak, P. (2002b). Talking to ourselves: The intelligibility of inner speech. *Behavioral and Brain Sciences* 25: 43-44.

Slobin, D. (1996). From "thought and language" to "thinking for speaking". In J. J. Gumperz and S. C. Levinson (eds.), *Rethinking linguistic relativity*, pp. 70-96. Cambridge: Cambridge University Press.

Slobin, D. (2002). What language is "mentalese"? *Behavioral and Brain Sciences* 25: 44-45.

Smith, M. (2000). Conceptual structures in language production. In L. R. Wheeldon (ed.), *Aspects of Language Production*, pp. 331-374. Psychology Press.

Smith, S. M., Ward, T. B., and Finke, R. A. (eds.). (1995). *The Creative Cognition Approach*. Cambridge, Mass.: MIT Press.

Smith, S. M. (1995). Fixation, Incubation, and Insight in Memory and Creative Thinking. In S. M. Smith, T. B. Ward, and R. A. Finke (eds.), *The Creative Cognition Approach*, pp. 135-156. Cambridge, Mass.: MIT Press.

Soto, D., Mäntylä, T., and Silvanto, J. (2011). Working memory without consciousness. *Current Biology* 21(22): R912–R913.

Soto, D. and Silvanto, J. (2014). Reappraising the relationship between working memory and conscious awareness. *Trends in Cognitive Sciences* 18(10): 520-525.

Sperber, D. and Wilson, D. (1995). *Relevance*. Second Edition. London: Blackwell.

Sridhar, S. N. (1990). As cited in Haspelmath 2007. *Kannada*. Descriptive Grammars. London: Routledge.

Swanton, J. R. (1911). Haida. In F. Boas (ed.), *Handbook of American Indian Languages*. Bureau of American Ethnology 40(1), pp. 205-82. Washington: Smithsonian Institution.

Sweetser, E. (1990). *From Etymology to Pragmatics: Metaphorical and Cultural Aspects of Semantic Structure*. Cambridge: Cambridge University Press.

Subramaniam, K., Kounios, J., Parrish, T. B., and Jung-Beeman, M. (2009). A Brain Mechanism for Facilitation of Insight by Positive Affect. *Journal of Cognitive Neuroscience* 21(3): 415-432.

Tamura, Y. (2014). Epistemicity and Deixis: Perspectives from Central Alaskan Yup'ik. *Berkeley Linguistics Society* 38: 494-506.

Tamura, Y. (2017). A Supplementary Description of the Use of Demonstratives in Central Alaskan Yup'ik. *Cognitive and Functional Linguistics II* [認知・機能言語学研究 II] 1(1): 31-40.

Talmy, L. (2000a). *Toward a Cognitive Semantics, Volume I: Concept Structuring Systems*. Cambridge, Mass.: MIT Press.

Talmy, L. (2000b). *Toward a Cognitive Semantics, Volume II: Typology and Process in Concept Structuring*. Cambridge, Mass.: MIT Press.

Taylor, J. (2012). *The Mental Corpus.* Oxford: Oxford University Press.

Thagard, P. (2007). Abductive Inference: From Philosophical Analysis to Neural Mechanisms. In A. Feeney and E. Heit (eds.), *Inductive Reasoning*, pp. 226-247. Cambridge: Cambridge University Press.

Thompson, C. and Kielar, A. (2014). Neural Bases of Sentence Processing: Evidence from Neurolinguistic and Neuroimaging Studies. In M. Goldrick, V. Ferreira, and M. Miozzo (eds.), *The Oxford Handbook of Language Production*, pp. 47-69. Oxford: Oxford University Press.

Tomlin, R. S. (1997). Mapping conceptual representations into linguistic representations: the role of attention in grammar. In J. Nuyts and E. Pederson (eds.), *Language and conceptualization,* pp. 162-189. Cambridge: Cambridge University Press.

Trübutschek, D., Marti, S., Ojeda, A., King., J.-R., Mi, Y., Tsodyks, M., and Dehaene, S. (2017). A theory of working memory without consciousness or sustained activity. *eLife* 6: 1-29.

Trübutschek, D., Marti, S., and Dehaene, S. (2019a). Temporal-order information can be maintained in non-conscious working memory. *Scientific Reports* 9(6484): 1-10.

Trübutschek, D., Marti, S., Ueberschär, H., and Dehaene, S. (2019b). Probing the limits of activity-silent non-conscious working memory. *Proceedings of the National Academy of Sciences* 116(28): 14358-14367.

Tuggy, D. (1993). Ambiguity, polysemy, and vagueness. *Cognitive Linguistics* 4(3): 273-290.

Turner, M. (1998). Figure. In C. Cacciari, R. Gibbs, A. Katz, and M. Turner, *Figurative Language and Thought*, pp. 44-87. Oxford: Oxford University Press.

Vicente, A. and Martínez-Manrique, F. (2005). Semantic Underdetermination and the Cognitive Uses of Language. *Mind and Language* 20: 537-58.

Vicente, A. and Martínez-Manrique, F. (2008). Thought, Language, and the Argument from Explicitness. *Metaphilosophy* 39: 381-401.

Vicente, A. and Jorba, M. (2019). The Linguistic Determination of Conscious Thought Contents. *No s* 53(3): 737-759.

Vigliocco, G., Kousta, S.-T., Della Rosa, P. A., Vinson, D. P., Tettamanti, M., Devlin, J. T., and Cappa, S. F. (2014). The Neural Representation of Abstract Words: The Role of Emotion. *Cerebral Cortex* 24: 1767-1777.

Voegelin, C. F. (1935). *Tübatulabal Grammar*. University of California Publications in American Archeology and Ethnology 34(2). Berkeley and Los Angeles: University of California.

Wallas, G. (1926). *The art of thought*. London: J. Cape.

Wang, J., Conder, J. A., Blitzer, D. N., and Shinkareva, S. V. (2010). Neural representation of abstract and concrete concepts: a meta-analysis of neuroimaging studies. *Human Brain Mapping* 31: 1459-1468.

Watkins, L. J. (1984). *A grammar of Kiowa*. Lincoln, Nebr.: University of Nebraska Press.

Wheeler, K. B. and Bergen, B. K. (2010). Meaning in the Palm of Your Hand. In S. Rice and J. Newman (eds.), *Empirical and Experimental Methods in Cognitive/Functional Research*, pp. 303-316. Stanford, Calif: CSLI Publications.

Wilson, D. and Sperber, D. (2002). Relevance Theory. *UCL Working Papers in Linguistics* 14: 249-287.

Wilson, N. L., and Gibbs, R. W. (2007). Real and imagined body movement primes metaphor comprehension. *Cognitive Science* 31: 721-731.

Yamanashi, M. (2000). Negative Inference, Space Construal, and Grammaticalization. In L. Horn and Y. Kato (eds.), *Negation and Polarity*, pp. 243-254. Oxford: Oxford University Press.

Yu, S. and Zenker, F. (2018). Peirce Knew Why Abduction Isn't IBE—A Scheme and Critical Questions for Abductive Argument. *Argumentation* 32: 569-587.

Yuan, Y. and Shen, W. (2016). Commentary: Incubation and Intuition in Creative Problem Solving. *Frontiers in Psychology* 7(1807): 1-9.