



Proprietary linguistic meaning

Tim Pritchard¹ 

Received: 30 December 2021 / Accepted: 9 June 2022
© The Author(s) 2022

1 Introduction

I wish to argue that a word's linguistically encoded meaning (the context-independent meaning of a word type) has, for very many words, a specific character and that this provides a way to distinguish a linguistically influenced conceptual grasp from more general conceptual content. To summarize my positive claim, what I call the proprietary nature of linguistic meaning arises due to a tendency, prompted by the use of a linguistic label, to seek commonalities across instances. This tendency reveals itself by a shift in focus away from specific attributes and towards shared relational characteristics of situations or items. This gives rise to a type of conceptual content that can be said to be encoded by a linguistic term, and which can be distinguished from conceptual content that is inferred from specific uses of language.

For example, consider the word 'bridge'. On the view that I argue for, the linguistic meaning of this word specifically relates to the characteristic of being an item that provides a pathway between two detached areas. The focus is not on the attributes of particular bridges, but on the fundamental nature of the relation to a situation that bridges manifest. As such, we can distinguish this specific and linguistically encoded meaning from more general and variable conceptual content pertaining to bridges, such as that they are frequently used to carry people and vehicles across valleys, rivers, and urban barriers; they are often made of steel, stone, and concrete; they can be flimsy (rope bridges) and formed from natural features (fallen trees); and so on. These other characteristics may be inferred from a use of 'bridge', but, by hypothesis, this will arise from accessing our general conceptual knowledge and situational specific knowledge, rather than being directly given by the language faculty itself.

Not all types of word provide clear initial motivation for this kind of hypothesis, and research has not been evenly balanced. Commonly discussed types of word in

✉ Tim Pritchard
t.pritchard@ucl.ac.uk

¹ Department of Linguistics, University College London, London, UK

philosophy, such as proper names and natural kind terms, may not be good exemplars. In psychology, Asmuth & Gentner (2017) note that empirical research has often ‘largely ignored’ the types of word that would most readily support this approach, in spite of the pervasiveness of such words. This, arguably, has led to the particular characteristic that I focus on being under appreciated.

I begin by discussing some recent views on word meaning in order to provide a context in which the hypothesis can be set (Sect. 2). I support my positive proposal by bringing in evidence from psychology (Sect. 3), morphology and syntax (Sect. 4), and philosophical reflection on the nature of understanding (Sect. 5). I propose that the evidence is strong enough to support the claim that there is a central tendency in language in how meanings are aligned with words, and that this plays a fundamental role in explaining our shared understanding of a language (Sect. 6).

2 Some recent views on word meaning:

‘carnivore’ [linguistic material]...  CONCEPT (of carnivores)

The left-hand side contains whatever belongs specifically to the language faculty. For ‘carnivore’ this could include, for example, a phonological form, morphosyntactic features such as singular/plural, the bound morphemes ‘carn’ and ‘vore’, perhaps an N categorizing head at some level of syntactic and morphological analysis.¹ On the views that I illustrate below, the core conceptual meaning of ‘carnivore’ is not specific to the language faculty but belongs to our extra-linguistic conceptual repertoire. The linguistic meaning of ‘carnivore’ is best described on these views as a pointer to (or point of access for, or instruction to fetch) an extra-linguistic concept. I don’t reject this view completely, but I will argue that the conceptual element is distributed across both sides of the arrow.

The notion of a concept of something is open to two different perspectives (Machery, 2009; Löhr, 2020). One perspective has a broad focus, looking at the kind of knowledge structures that are studied in psychology and that are required to explain cognitive abilities such as our ability to identify (e.g. that the animal standing in front of us is a carnivore) or draw inferences (e.g. that it probably has sharp teeth). Related abilities include typicality judgements and the ability to list attributes.² A second perspective has a narrower focus, looking at what is involved in thinking of something as a such-and-such. When we identify something as a carnivore and make inferences about it, we are also thinking of that thing as a carnivore. A person who cannot identify the animal as a carnivore, or who does not draw all the same inferences, can still understand what it is for something to be a carnivore. While an explanation of the two perspectives may overlap, the perspectives themselves are different. I will comment later on this distinction (Sect. 5) as it is relevant for my claim that linguistic

¹ See Embick (2015) for categorizing heads, a feature of a ‘distributed’ approach to morphology.

² Hampton (2015) surveys prototype and exemplar theories of concepts in relation to discussions of word meaning.

conceptual material has its own proprietary nature. But for now I will borrow from Glanzberg's neutral description of concepts as 'units that are active in thinking, and contribute content to cognitive states' (Glanzberg, 2018, p. 202).

Glanzberg's (2011; 2014; 2018) account of word meaning illustrates the pointer/concept framework. In his account, non-functional vocabulary items, such as nouns, verbs, and adjectives, have root meanings that are said to come from extra-linguistic conceptual material. For example, a word such as 'open' is associated with a root meaning OPEN which expresses the idiosyncratic notion of being open (whatever it is that is specific to being open).³ The strictly linguistic side of 'open' acts as a pointer to the concept that occurs outside the language faculty (2011, p. 15). Importantly, any use of 'open' is also packaged by linguistically provided material that provides 'a kind of structural frame' (2014, p. 28). So, while the root OPEN will correspond to an extra-linguistic concept, the meaning expressed by a use of the word 'open' will also be determined by the linguistic packaging that accompanies the use of the word. For example, the use of 'open' in 'Helen opens the door' is accompanied by a structural frame that determines the kind of event that is being described:

[[[x ACT] CAUSE [BECOME [y ⟨OPEN⟩]]]]

Our interpretation of the sentence is informed by an event structure, such that we treat Helen as an agent who causes the window to take on the property of being open. This event structure is found in many sentences ('Helen shuts the window', 'John paints the door'), and is encoded as part of the linguistic system itself (whether lexically or syntactically). This contrasts with the specific idiosyncratic meaning of 'open', to which the word points but which is in fact part of our extra-linguistic cognitive repertoire. Glanzberg describes the root (OPEN) as an interface between the language faculty and our wider cognitive repertoire (2011, p. 15). The conceptual material corresponding to a root invariably occurs within a linguistic packaging, which can consist of overtly expressed functional elements such as determiners, quantifiers, inflectional and derivational morphology, and covert syntactic material that, for example, on some views conveys event structure (e.g. Ramchand, 2008) or category defining functional heads (see Embick, 2015, p. 43).⁴

Pietroski (2018) likewise distinguishes the notion of linguistic meaning from conceptual material. Pietroski is explicit that he wishes to explain linguistic meaning not in terms of a concept but in terms of an instruction to fetch a concept. The meaning of 'cow' can be represented as the instruction *fetch@'cow'*, where for each lexical item L, the instruction *fetch@L* is executed by copying a concept that resides at the long-term address of L into the workspace (2018, p. 106). The instruction is distinct from the concept that resides at the address (or, rather, family of concepts, so as to

³ This is a semantic, not a morphological, notion of a root: 'root' refers to the specific idiosyncratic concept that a word conveys and that can be distinguished from more general functional contributions to meaning. See e.g. Levin & Rappaport Hovav (2005, p. 71). How this conceptual material relates to particular levels of morphological analysis of a word is debated.

⁴ In line with these ideas, and following Acquaviva (2019), I take it that denotation and predication only arise within larger syntactic structures (e.g. nouns appear in a DP), so it may not be appropriate to style the linguistic meaning of an N or V as type <e t>.

deal with polysemy; see 2018, pp. 4–6). Linguistic meanings are ‘concept assembly instructions’ (2018, p. 292), and the meaning of a syntactically arranged string of words ‘is a complex instruction’ (2018, p. 294). Whereas Glanzberg makes a distinction between the functional elements of language, which directly encode meanings, and lexical categories such as nouns/verbs/adjectives, which act as interfaces to extra-linguistic conceptual material, Pietroski has a more uniform approach in which even functional elements are treated as instructions to access an appropriate concept.

Carston (2016; 2019) outlines a view that has some similarities to Glanzberg and Pietroski. She distinguishes a broad communicational perspective on the lexicon (the C-lexicon), which stores the senses that are expressed in actual communication, from a narrow perspective (the L-lexicon) that considers merely what is necessary for specifying language considered as an internal computational system. From the L-lexicon perspective, there is debate over whether the components of the lexicon are to be treated as categorized units, such as nouns and verbs, or categoryless roots (here with the morphological notion of a root in mind, viz. a recognizable recurrent form in a word that can undergo affixation and other morphological operations). In particular with polysemy in mind, Carston (2019) notes that there is also debate on what notion of meaning is relevant for these units. The two options she mentions are an underspecified ‘core’ meaning that is common to all the lexicalisations based on a given root, or, corresponding to Glanzberg’s and Pietroski’s accounts, merely an index or address. Polysemy is also a motivating factor for Pietroski. If a word can be used to express a family of related concepts, we can’t identify any one concept as the particular meaning of a word. The linguistic meaning of a word (the *fetch@L* instruction) will fetch the family of concepts, and thereby acts as a constraint on expressed meaning without fully determining what is expressed on an occasion.⁵

The key point for my purposes is that these accounts separate a specifically linguistic side to meaning from a more general notion of conceptual content, which the terms we use can point us towards. Glanzberg describes his account as offering ‘a clean division of labor between distinctively linguistic aspects of cognition and our wider cognitive abilities’ (2018, p. 14). I am entirely sympathetic with a division of labour, but on the view I support, the word ‘carnivore’ will immediately encode not merely for phonological and morphosyntactic features but also for a specific type of conceptual content that is specially enabled by linguistic form itself: ‘carnivore’ will linguistically encode for a shared relational characteristic that holds of items denoted by ‘carnivore’, namely, an animal that feeds on flesh. Use of the word may point us to further elements in our general knowledge of carnivores, such as typical appearances or behaviours, and inferences drawn from specific situational knowledge—precisely the kind of information that the study of our concepts of carnivores will typically

⁵ Polysemy is a difficult challenge for any account of word meaning, and I don’t attempt to address it in this paper. However many (related) meanings a word may have, my aim is to clarify a central feature of those meanings. Suffice to say that the account I propose stands in tension with both the ‘pointer’ response to polysemy (also see Quilty-Dunn, 2021), and a ‘semantic feature’ hypothesis proposed in Brocher et al., (2018). Quilty-Dunn construes concepts themselves as pointers, and by aligning concepts with linguistic meaning gives us another example of a view that corresponds to the ‘pointer’ pattern that I am illustrating here.

consider. But this further information, by hypothesis, is not part of our linguistic knowledge.⁶

There are related views of word meaning that can be interpreted as, in effect, putting more material on the left-hand side of the pointer. These are views that look to some sort of conceptual commonality that is said to hold across different uses of a word (Ruhl, 1989; for discussion and criticism, see Carston, 2016, Sect. 2.2; 2019; Recanati, 2017; Vicente, 2018). The suggestion is that we can abstract from what is specific to individual situations and end up with a common core that occurs across different uses of a word. For example, an open bottle, and opening a bottle, differ in specific respects from an open window and opening a window. But it is plausible to think that there is also something in common across both situations—something like an item being in (or put into) a state in which it permits access to some space that by design or nature it otherwise prevents access to.

In some respects, this is a completely different account to the pointer/instruction account: a representation of some core feature of a concept is not an instruction to fetch a concept. To illustrate with an analogy taken from Eliasmith (2013), whereas the address of an image that we are instructed to fetch can be represented simply by a number (e.g. fetch image 32), a compressed version of the image will still be an image as such. We are used to this from JPEG files: the picture we see on our computer screen will typically be heavily compressed in comparison to the original photograph. The compressed version is not as such an instruction to fetch the full version, nor is it a label that shows where the full version is stored.

On the other hand, if we think of the interpretations and inferences that arise in communication, it is clear that we will often go beyond any putative common core conceptual content. Suppose I say ‘The door is blocked’. By hypothesis, our grasp of English only provides us with a general understanding; we will need to know more about the specifics of the situation before we can draw practical conclusions (what door are we talking about; what sort of blockage; what are our purposes; and so on).

There is nothing new about proposing linguistically influenced distinctions within the general notion of a concept. Many, though not all,⁷ are sympathetic to the idea that there is probably a distinction to be made between lexical knowledge (the type of knowledge that is required for having a bona fide proper understanding of a word) and more general encyclopaedic knowledge. The difficulty is in finding a principled way to make this distinction. The evidence that I discuss in this paper provides one way of motivating the claim that there is a distinction.

⁶ Here I am echoing Bierwisch’s ‘two level semantics’, in which he distinguishes a ‘semantic form’ (that comes from knowledge of language per se), and ‘conceptual structure’ (which corresponds to the situational and specific understanding that may arise from an utterance). As Bierwisch and Schreuder put it: ‘the conceptual structure, in terms of which the actual interpretation of linguistic expressions is specified, merges the conditions specified by semantic form with information coming from different domains of encyclopedic background knowledge, contextual information and situational conditions’ (Bierwisch & Schreuder, 1992, p. 32).

⁷ Peeters (2000) reviews several contributors to the debate, and notes that representatives of Cognitive Linguistics in particular have denied the distinction. Peeters cites Haiman (1980) as an important influence on this denial of the distinction.

3 Evidence from psychology: linguistic labels and thinking analogically

Work on analogical cognition by D. Gentner and her colleagues provides evidence that the use of linguistic labels promotes a specific kind of cognition.⁸ Analogical reasoning is the ability, highly developed in humans, to think of situations and items in terms of shared relational characteristics rather than just in terms of the particular attributes of specific instances. Take the example of bridges once again. We can identify the particular attributes of specific bridges (such as their colour or size) and attributes that are often shared (many older bridges are made of stone, many newer ones are made of steel; many take road traffic, others take rail traffic; and so on). A relational focus, by contrast, considers relations between items. For example, a bridge is understood (roughly speaking) to provide a pathway across a gap. Here we are not describing the attributes of a given bridge or group of bridges. Instead, we are thinking of a bridge qua something that plays a particular role in relation to a gap. By virtue of this ability to look beyond specific attributes to more abstract relational characteristics, we have no difficulty in grouping together items that may differ visually in almost every respect (contrast a stone bridge across a stream with a modern road bridge). We are able to group items in terms of how they relate to the situation they occur in rather than how they appear.

A simple test for analogical reasoning ability is the relational match-to-sample test. This test illustrates the ability to match items on the basis of a relation that they manifest. For example, AA matches to XX not to CD because AA and XX are identical in that they both manifest an identity relation between their constituent members. Christie & Gentner (2014) used a version of this test to show that infants who otherwise fail the test are enabled to succeed when the task is aligned with a linguistic label. Their version of the test used samples consisting of three pairs of coloured shapes (such as two green ellipses, two purple rectangles, and two orange non-matching shapes). The children were shown a sample (the standard) and then asked which of two other samples matches the sample. One of those other samples was a relational match (the two green ellipses go with the two purple rectangles).

Christie and Gentner tested 2- to 4-year olds. The 2- and 3-year olds did not show an ability to make the matches on a relational basis. To confirm that all the children realized that pure relational similarity counts as a match, a second experiment was conducted in which the children received initial training and corrective feedback on their performance (confirming a correct choice of match or showing the child what the correct choice was). The initial experiment was then re-run, but the 2- and 3-year olds still did not show an ability to make the match on the relational basis. In a third experiment, the children were given initial symbolic training with the terms ‘same’ and ‘different’. The initial experiment was re-run (with no use made of ‘same/different’), and here the results showed that 3-year olds were helped to succeed in the task.

But the most striking results came from a fourth experiment. In this experiment there was no initial training or feedback. The experiment was the same as the initial experiment except for one difference: ‘The only difference was that a novel label

⁸ See Pritchard (2019) for a review of the research as it relates to word meaning.

(“truffet”) was applied to the standard, and the child was asked which of the two alternatives was also a truffet’ (2014, p. 391). In this condition, all age groups chose the correct relational match at above chance levels. The authors comment:

This language benefit is quite remarkable considering that the label applied to the standard was a novel word, and that children received no training or practice trials. These results are consistent with prior findings indicating that application of common words ... results in relational learning. (2014, p. 392).

These findings are in part explained by the hypothesis that use of a label invites a search for sameness across items or situations. The search for sameness involves making comparisons, and the result of making comparisons while looking for a commonality is a shift in focus away from the highly variable specific attributes of individual items and towards shared relational characteristics. This impact of the search for a comparison (a shift to relational characteristics) is widely illustrated and tested for in a considerable body of experimental literature, and is one of Gentner’s key findings (for overviews, see Gentner & Smith, 2013; Gentner & Hoyos, 2017). The presence of this shift in our response to individual words is well motivated. A single label is a serious prompt to look for an expected sameness of some sort across the different instances named by the label. This will not in general be provided by comparing individual attributes of an item. Namy & Gentner (2002) found that giving children a single label for two members of an object category encouraged a shift towards relational alignment of the members, whereas conflicting labels (on the same task) undermined this shift. Similar alignments between the use of labels and a relational orientation were found in Gentner et al., (2011) for 3- to 6-year olds, and Goldwater & Markman (2011) for adults.

This work on analogical cognition motivates the hypothesis that the use a word orients us towards a relational focus. The claim is not that a linguistic label is essential for this focus, but rather that use of a word (or at least of many words) prompts and helps to sustain this perspective (see Gentner, 2016).

I want to distinguish this claim from other work in psychology that has also looked for a special influence from linguistic labels. In a series of papers, Lupyan has argued that linguistic labels ‘do not simply refer to nonlinguistic concepts, but rather actively modulate object representations’ (Lupyan, 2012, p. 255). In other words, the type of mental representation that use of a label induces is claimed to have specific features that highlight particular aspects of the represented objects. Rather than merely being passive pointers to a conceptual representation, the use of a word influences the form of the representation itself. Conceptual representations activated by words are different from those activated by nonverbal means.

Lupyan in particular suggests that use of a linguistic label leads to our mental representations becoming more ‘categorical’ in nature:

Simply calling something by its name may shift the representation of the labelled object such that properties typical or diagnostic of the category are highlighted while properties irrelevant to the category are abstracted over. (Lupyan, 2012, p. 268).

These are features that enable us to recognize that two or more objects/events belong to the same category, with categorization described as ‘the process by which an input is aligned in some way with that of previously encountered members of same category’ (Perry & Lupyan, 2017, p. 925). This ability is enhanced by the linguistic label leading to a focus on diagnostic properties. For example, colour is not an informative feature for whether something is a cup, so the mental representation that is activated when we use the word ‘cup’ will not highlight that feature. This choice of particular features is called ‘selection’, which ‘can be thought of ... as a warping of a representation ... into a task-relevant form’ (Perry & Lupyan, 2017, p. 925).

Evidence for this comes from studies that show that named categories are more easily learned than unnamed categories. In the ‘grecois aliens’ study (summarized in Lupyan 2012, pp. 265 f.), visual representations of two types of alien (dangerous and not dangerous) were presented to participants. In an initial learning phase, participants familiarized themselves with the two representations and the type of alien they depicted. The experimental task was to respond with the appropriate direction of motion (escape/approach) to a picture of an alien. One group of participants was also given a linguistic label for each type of alien, while the second group of participants was not. The results show that participants who learned the category distinction along with a label learned to recognize each type of alien about twice as fast as those in the no-label condition.

The explanation for this effect is not dissimilar to that proposed by Gentner. Lupyan and his colleagues suggest that the use of a linguistic label ‘promotes abstraction’ (Lupyan & Lewis, 2019, p. 9). While a perceptual cue of an item will be accompanied by specific characteristics of that item, use of a word for an item does not. This leads to the representations associated with words abstracting away from the idiosyncratic information that might happen to characterise a particular instance of an item (Lupyan & Lewis, 2019, p. 10). Experimental results show that where labels are used, people have poorer memory for the idiosyncratic features of test items (Lupyan, 2012).

Lupyan’s results provide interesting evidence that use of a linguistic label has some sort of biasing effect on our mental representations. Rather than just being like a label on a box file that contains general conceptual information, a linguistic label leads to some kind of selection in which a representation is ‘warped’ into a task appropriate form.

On the other hand, Lupyan’s focus is on recognition of instances of a category, and on making the kind of inferences from a category member that we typically can make on seeing an instance. There are two potential issues with this if we are focusing on word meaning.

First, our use of words is not geared specifically to categorization, where that is construed as an ability to recognize instances. When we process texts or speech, we may only rarely (or not at all) need to actually connect parts of what is said with specific features of the environment around us. We need to understand the ideas that are expressed, but it does not follow that we will have an ability to categorize. You may tell me that there is explosive in the room and I can process and understand that utterance, without my being able to indicate which thing in the room happens to be the explosive. Conversely, a trained dog may be able to pick out the explosive but

clearly without any understanding of the linguistic term ‘explosive’. In general, it seems that the relevant ability that is supported by the kind of mental representation that is specifically associated with a word is not the ability to recognize instances. Understanding a word in an utterance, at least in some initial basic way, does not depend on categorization.

Second, many of the inferences that we typically make arguably cannot arise simply from knowledge of a word per se. We need to complement this with general background knowledge and situation specific knowledge. If you are driving along a road and I tell you that there is an obstacle ahead of us, no particular inferences can be drawn until you see the position, size, and nature of the obstacle, along with the nature of the road at that point. Our linguistic grasp of ‘obstacle’, ‘road’, ‘ahead’, doesn’t provide us with the relevant information. Of course, you will naturally slow down, or stop, even without any further information, but this will arise from the general recognition that people give warnings like this when they potentially effect a current activity.

Lupyan’s results are interesting and are not inconsistent with the orientation towards relational commonalities that Gentner’s work supports. But categorizational and inferential ability is not the right account of the specific kind of cognitive achievement that we attain to in our grasp of a word meaning. In Sect. 5 I describe how work in philosophy connects grasp of structure with ‘understanding’, and I will suggest that this in itself (an understanding) is the appropriate focus that a tendency to look to relational commonalities gives rise to. Before this, in Sect. 4, I describe some evident ways in which our knowledge of language includes an abstract relational element.

4 Evidence from morphology and syntax

The suggestion that linguistic knowledge has an orientation to abstract relational commonalities seems to be well supported by the kind of understanding that evidently informs our grasp of morphological and syntactic patterns.

For example, the use of the -er suffix in English very often aligns with an interpretation in which the external argument⁹ of a verb is picked out (see Rappaport Hovav & Levin, 1992): walker (the person who walks), sprinkler (the item that sprinkles), compiler (that which compiles), twister (that which twists), and so on. The data is complex and not all uses of -er fit the external argument (Booij & Lieber, 2004), but this generalization clearly describes a major component in our understanding of this suffix. As such, our understanding of what this suffix contributes is highly abstract and relational: it indicates (in many uses) that we are to think of whatever it is that corresponds to a particular syntactic position. This is even more abstract than the various thematic roles that can be exemplified by the external argument (agent, instrument, theme, source; see Rappaport Hovav & Levin, 1992).

⁹ The external argument corresponds to the subject position in a sentence, being distinguished from the internal argument which corresponds to the complement position. ‘The boy gave Mary the book’: ‘The boy’ is the external argument, ‘Mary’ and ‘the book’ are the internal arguments. Arguments can be associated with various thematic roles: ‘The boy’ is agent; ‘Mary’ is goal; ‘the book’ is theme.

We only need to peruse a list of common derivational affixes to gain a clear sense of the many different examples of relational understanding that inform our grasp of language. The following examples are taken from Katamba and Stonham (1993/2006, pp. 50 f.), along with the suggested meanings of the affixes.

Prefixes:

- re- (again): rewrite.
- ex- (former): ex-mayor.
- en- (put in): en-cage.

Suffixes:

- -ment (result of doing the action): arrangement.
- -less (without): powerless.
- -ful (having): powerful.

The accuracy of the glosses on the meaning is not the issue here. The important point is the evident way in which all of these (and many other examples can be given) pick up on relational notions. A re-write is understood to be a writing that has a particular relation to a previous writing; an ex-mayor is a person who has a particular relation to the office of mayor; to be powerless is to stand in a particular relation to the having of power; an arrangement is the result of an act of arranging.

We can also look to the kind of understanding that is implicit in our grasp of syntax. Notions of tense (future, present, past), aspect (punctual, progressive), person (speaker, addressee, ...), number (singular, dual, plural ...) all rest on our ability to grasp highly relational notions. Some theorists align syntax with event-structure, in which it is taken for granted that language users comprehend abstract and relational notions such as ‘cause’, ‘become’, ‘result’ (e.g. Ramchand, 2008). Notice that if we are asked to group together a series of causes, we do so not by identifying intrinsic features of a particular item or event but instead by identifying when the required relation holds with an effect.¹⁰

Interpretation of compound words also illustrates the importance of our ability to grasp abstract relations (see Gagné & Spalding, 2013; Olsen, 2012). Our comprehension of words such as snowman, handbag, trumpet blast, health worker, news programme, mountain stream, and so on, depends on our ability to grasp a number of relations that can hold between the two constituents (such as MADE OF, LOCATION, SOURCE ...).

In other words, our understanding of a language is pervasively marked by an understanding of highly abstract relational patterns. This is necessary for us to understand syntactic constructions and morphologically complex words. In itself this is not enough to show there is a general tendency towards a relational perspective in our grasp of word meanings. Indeed, it is common to make a distinction between

¹⁰ Glanzberg (2011, pp. 19–21) has a discussion of (syntactic) CAUSE and indicates how it seems to differ from our ordinary notion of cause. But even so, it still has the same general abstract nature—indeed, as Glanzberg notes, more abstract than our ordinary notion of cause.

functional vocabulary (and other elements that provide the structural packaging that Glanzberg describes), and open class terms (such as nouns and verbs). But the following points can be made.

First, if the fundamental syntactic and morphological structure of a language is built on our understanding of relational notions, it would seem at least *prima facie* plausible that this will also influence our understanding of open class terms. In both cases, we have linguistic units that apply across multiple instances, and as we have seen there is good evidence from experimental psychology that this enhances a relational perspective. Further, if we share an understanding of the words we use, this motivates the hypothesis that words will, as it were, tend to shed requirements on understanding that won't in general be shared between users. Knowledge of particular attributes of bridges and carnivores will not in general be shared; if we genuinely share a basic understanding of the words we use, we can explain this if the language-specific requirement on understanding is given just by the relational characteristics (item that provides a path across a gap; animal that feeds on flesh).¹¹

Second, there is no absolute barrier between bound morphemes and free-standing words. Fabregas & Scalise (2012, p. 9) discuss the Spanish affix *-ismo*, which can be paralleled with the English affix *-ism*. Usually this appears as a bound form (Marxism, surrealism, masochism), but it can also appear as a free form: 'They do not defend any liberty or attack any ism'; 'She certainly avoided any kind of ism in her writing' (examples from the *Corpus of Contemporary American English*). This may suggest that there is no fundamental distinction between functional vocabulary elements and the kinds of word that appear within the structures that the functional vocabulary helps to establish.¹²

Third, inspection provides strong intuitive support for the claim that a relational factor plays a major role in our grasp of many words.¹³ Consider: cause, effect, solution, problem, indication, episode, vengeance, sequel, progress, failure, deficiency, predator, prey, pest, harvest, promotion, filler, answer, weapon, gift, omnivore, and so on (we could simply pick out most of the nouns and verbs used in this paper). Relational characteristics are evident throughout: a filler is something that is used to fill a space, an answer is something that stands in an appropriate relation to a question, a weapon is something used to cause harm and injury, a gift is something that is given in a particular way, an omnivore is something that feeds on both plant and animal matter. Asmuth & Gentner (2005) estimated that half the nouns in a representative

¹¹ I return to this point in Sect. 6. Implicitly, I am taking a position in which the 'semantic facts' about a word are grounded in facts about speakers' behaviour (see Fricker, 1982), though giving this a cognitive orientation, with a claim about how a particular conceptual focus naturally arises from the use of a linguistic term by a language community, and thinking of the semantics of a word as whatever it is that our understanding grasps by virtue of being competent users of a language.

¹² I am giving the phonological exponent of some morphological entity a key role in explaining the type of meaning that a linguistic term expresses. This may not be consistent with the kind of architecture found in distributed morphology, in which syntactic processes act as an input to two separate interfaces, phonological and semantic (see Embick, 2015). On the other hand, I am not making any explicit claim about the meaning of a morphological root. Acquaviva (2014) argues that roots 'are not the formal side of a sign whose content is the core lexical meaning'.

¹³ This point is much more general than the relatively limited notion of 'relational terms' that have been studied in linguistic semantics. See Pritchard (2019, pp. 589–590).

corpus of adult vocabulary have this characteristic. And yet, curiously, such words have often been overlooked, with Asmuth & Gentner (2017) reporting that they have until recently been ‘largely ignored’ in the empirical psychological literature. And while there are hundreds of philosophical papers on demonstratives, proper names, and natural kind terms, philosophers likewise have mostly ignored this pervasively common type of word.

A fourth motivation for aligning linguistic meaning with relational structures comes from considering understanding.

5 Evidence from philosophy: the nature of understanding

Understanding appears to have a distinct role to play in our epistemic life, having characteristics that arguably distinguish it from knowledge as such (see Baumberger et al., 2016).¹⁴ Amongst philosophers working on understanding, there is widespread agreement on the following intuition: understanding involves ‘the discernment of significant structure’, seeing something ‘either as having an internal structure or as an element in a larger pattern’ (Franklin, 1981, p. 203). Zagzebski speaks of ‘grasping relations of parts to other parts and perhaps the relations of parts to a whole’ (2001, p. 241). Understanding is ‘the grasp of order, pattern, and how things hang together’ (Riggs, 2003, p. 217); ‘an internal grasping or appreciation of how the various elements in a body of information are related to each other’ (Kvanvig, 2003). Theorists typically talk of grasping or discerning these relations, in order to suggest the kind of knowledge that comes from a direct awareness of the relations involved (see Bourget, 2017).

For example, one type of understanding of the London Underground may be shown in the ability to navigate the system. This involves, *inter alia*, grasping the connections between the lines, how they relate at the various stations, and how the trains move between the stations. This understanding mainly involves a grasp of the internal structure of the Underground. Another type of understanding would be shown in an awareness of why the London Underground was built in the way it was. This will in part involve knowledge of how the building of the Underground was influenced by external demands, such as the geography of London and the competition between different companies involved in building their own lines. This is an understanding that sees the connections between the Underground and the larger situation in which it is set.

Again, a person may have a lot of knowledge about a Bach Fugue, such as knowing all the notes by heart, time and place of composition, and so on. But understanding the Fugue could be said to depend on appreciating how the theme is expressed in several voices, along with grasp of how that theme is varied (augmented, inverted), and grasp of how the various keys into which the piece modulates relate to each other. Alternatively, understanding a Fugue may involve grasping how the Fugue relates to

¹⁴ Pettit (2002) argues that linguistic understanding and knowledge differ with respect to Gettier cases, warrant, and belief. For recent discussion (though focusing on understanding what is said) see Longworth (2018).

an accompanying Prelude, or to a previous Fugue in the same key, or how Bach may have developed the Fugue form in relation to what was usual at the time. In each case of understanding, we are considering either the significant relations that are exhibited within the structure of the Fugue, or how the Fugue relates to other external factors.¹⁵

A connection between understanding and seeing how things relate is suggestively indicated by Franklin's (1981) observation that while we may say 'I know when he arrived', we cannot ordinarily say 'I understand when he arrived'. He suggests this is because a time of arrival does not, as such, have the sort of structure that invites an understanding (cf. Cooper, 2000, p. 384, on understanding a stone). We may, of course, relate a particular time to other times and events, and thereby get an understanding of why a person arrived at the time they did; but this is not to understand when he arrived.¹⁶

While we do not ordinarily speak of understanding when someone arrives, we do often very naturally speak of understanding, or not understanding, a word. As well as speaking of knowing the meanings of words, we speak of understanding the meaning of words, or just of understanding the words. If theorists are correct that there is an interesting notion of understanding at play in our epistemic life, this evidence from the ease with which we use 'understanding' of words may indicate that we can draw on this notion when considering word meaning. This is supported by the claim that language use directs us towards relational commonalities, for this is precisely the discernment of significant structure. We find, therefore, a natural alignment between the philosophical characterization of understanding and a central feature in our use of words. This in turn is consistent with the centrality that understanding has in our response to language.

We can illustrate by considering the words 'target' and 'barrier'. Roughly, the word 'target' indicates something that we aim in some manner to reach, and 'barrier' indicates something that prevents movement. I take it that, in the normal case, a condition for understanding these words is that we grasp these characteristics and know that 'target' and 'barrier' express them. We are considering the relational structure into which possible denotations of 'target' and 'barrier' fit.

The target/barrier examples show the following characteristics. First, the expressed relational structures are simple and plausibly can be shared by people in general. I will call our understanding of a word a basic understanding, to distinguish it from the fuller understanding that we may typically possess about items and situations. Notice that understanding the words 'target' or 'barrier' is not to understand why some particular thing is our target or why some particular thing is actually a barrier for us. In particular contexts we might disagree about or not understand these things. But by separating out a basic understanding, we have a way of preserving the claim that speakers of English generally share, across contexts, an understanding of the words 'target' and 'barrier'.

¹⁵ For other examples, we can think of the distinction between understanding a person and knowing facts about a person, or between the understanding of a game of football that a manager may have and knowing facts about that game of football.

¹⁶ We do also say 'I understand he arrived at 4pm'. This, arguably, is a case of polysemy, in which 'understand' is used to mean something like awareness of or acknowledgement of.

Second, we can distinguish this grasp of a basic relation from the understanding that experts may possess about targets and barriers. There will be experts who know how to build barriers of various sorts, who know how to identify targets in armed conflicts, and so on. This expert understanding will have particular characteristics: being graded, it will vary between people; it is hard to attain so counts as a praiseworthy cognitive achievement; it will be manifested in special abilities to infer specifics about the characteristics that particular targets or barriers will or should have; there will be a high level of integration with many other areas of knowledge. By contrast, if grasping a basic relation provides the fundamental element of our understanding the words ‘target’ or ‘barrier’, a linguistic understanding will not share these broader characteristics that understanding in general displays.

Third, if understanding the words ‘target’ and ‘barrier’ involves grasp of a significant relation, a natural connection will hold between understanding the words and understanding what targets and barriers are. Often, understanding what something is can at least partially be satisfied by grasping how something fits into the world, and this I take to be answered by the general relation to surrounding things that is expressed in the proposed basic understanding of ‘target’ and ‘barrier’.

Finally, we can also connect this perspective with the view of concepts that looks to the requirements for thinking of something as a such-and-such (see Sect. 2). On this view, our understanding of the word ‘barrier’ provides the basic orientation towards an item that is presupposed when we talking about barriers.¹⁷ Further support for the importance, when studying language, of this distinction between two different views of concepts comes from Gagné & Spalding (2013), who look at our comprehension of modifier noun-phrases and compounds. They propose a distinction between what they call the initial stage of comprehension, in which we comprehend a ‘relation-based gist’, from a subsequent elaboration in which we derive particular implications based on our understanding of the relation and attributes of the coordinated items. For example, ‘snowman’ gives a relational-gist interpretation (man MADE OF snow), which may then be elaborated by the inference that this will be a cold item (Gagné & Spalding, 2013, p. 104). In particular, they expand on their results and make an intriguing suggestion (see pp. 121–124) based on experimental work on what is known as the modifier effect in property verification tasks. Processing the phrase ‘peeled apple’ makes it harder to affirm that apples are red (2013, p. 123, reporting results from Gagné & Spalding, 2007). Lack of redness is a natural inference to make for peeled apples, but if our apple concept included features such as redness we may wonder why verifying redness is hindered even though the word ‘apple’ occurs in the noun phrase. This effect is robust and there are competing explanations.¹⁸ Gagné and Spalding suggest that the results support a view of concepts in

¹⁷ Fodor (1998; 2008) treats a concept as that which explains our ability to think about an *x* as an *x*. He construes this in non-epistemic terms as possession of a mental particular that has an appropriate causal-nomological connection with *x* things. Fodor indicates that concept possession provides us with a ‘primitive’ way of thinking of something, and glosses this as the idea that we may think in ‘file names’ (2008, p. 95). This suggests an extremely minimal requirement on being a constituent of thought, at the limit being nothing more than a symbol in a language of thought. I am arguing for a richer view, though the file-name idea may be more relevant for some types of word.

¹⁸ See Gagné & Spalding (2014).

which features (attributes we may associate with items that fall under the concept) are not part of a concept per se, though may be inferred. Predicating a particular feature, such as that apples are often red, or that bridges are often made of stone or steel, 'is a separate cognitive act from those involved in creating the concept or in using the concept to represent the thing to the mind' (2013, p. 124).

As well as supporting the idea that we should distinguish a basic conceptual grasp from knowledge of features that are inferred, this work also suggests a method for empirically testing the claim that some conceptual material is encoded as part of linguistic meaning. If this claim is correct, we would expect to find (if Gagné and Spalding's reasoning is correct) differences in behavioural (or neuro-physiological) responses to property verification tasks, depending on whether the property we test belongs to the proposed relational core meaning or not.¹⁹

6 Discussion

The position I am opposing is the claim that, aside from some linguistically encoded structural elements of meaning, there is nothing specifically linguistic to say about the conceptual material that use of a word aligns with. Against this, I have suggested that the use of linguistic terms promotes a particular kind of cognitive orientation. There is empirical evidence that words act as attractors for a specific way of thinking. Use of a linguistic label prompts a search for sameness across instances, and this results in a focus on shared relational structures. The relatively abstract nature of this focus is amply displayed in the understanding that underpins our grasp of morphological and syntactic regularities. This orientation appears to fit well with independent reflection on the nature of understanding.

The proposal is that there is a general orientation in linguistic meaning away from specific attributes and towards shared relational structure, with this in turn underpinning a basic understanding. However, not all words provide clear support for this claim. I have considered words for items and situations that are common parts of our everyday experience (barrier, carnivore, bridge, etc.). But here I want to look at some types of word that may raise questions.

Many words in scientific contexts are used to name items that are utterly remote from ordinary experience, such as terms for different types of elementary particle (leptons, quarks, higgs bosons). At the level of the ordinary language user, understanding will be restricted to generalities such as 'type of elementary particle'. One option here is to propose that, for such words, the linguistic meaning encodes just for this shallow understanding. Given this, there could be some tendency to think of such words as pointers to, or placeholders for, an entity about which we otherwise have little understanding of (cf. Millikan, 2000, on substance terms).²⁰ This would be an

¹⁹ Hampton et al., (2011) find little evidence for an impact of property centrality on the modifier effect (except under one condition, in experiment 6). However, the 'central properties' they used (see their Appendix B) did not in general reflect the kind of relational property that I have argued for.

²⁰ Millikan typically concentrates on words for which a placeholder focus has prima facie plausibility. I do not think that this gives us a good basis for developing a general theory of linguistic meaning (see Pritchard, 2017).

apparent convergence on the pointer views outlined in Sect. 2, but even here some minimum of understanding will be lexically represented. If you don't know that a lepton is an elementary particle, arguably you do not have a proper understanding of this word. The point can be applied to terms for things that are less remote from our everyday experience, such as names for diseases, or types of tree, about which our understanding may go no further than that it is some disease or some type of tree (cf. Putnam, 1975, on elms).

There are also intermediate examples, such as 'carburetor'. The kind of relational factors that mark out carburetors are readily accessible (a device that mixes fuel and air for combustion), though also not known by perhaps the majority of language users. Here I suggest that we can leave it as an open question whether someone who only knows that a carburetor is part of an engine has an incomplete linguistic grasp of the word 'carburetor', or whether the linguistically encoded material is more limited.

A different type of question arises for words used for the complex biological and chemical kinds that surround us, such as 'cat', 'tree', 'silver'. In contrast with leptons and carburetors, ordinary users of English typically have quite detailed knowledge of specific attributes of cats and trees, and it may not initially seem appropriate to think in terms of relational characteristics. Yet even here, there is motivation to focus elsewhere than on specific attributes (such as colour and shape) when considering what 'cat' (for example) may linguistically encode. As well as Gagné and Spalding's suggestion mentioned above that specific attributes may not be part of a concept per se, suggestive evidence also comes from work on natural kind terms by Keil (1989, reported in Pritchard, 2019). This work indicates that even pre-school children can override surface perceptual features when making judgements about kind membership in some situations. In addition, relational factors can be argued to play a role here as well. A referee for this journal suggests that, for cats, their expected actions and relations to the environment could play a significant role in our understanding of cats. This idea is supported by Mandler:

animals are objects that move in different ways than artifacts do and are seen to move themselves rather than being moved by others. For both animals and artifacts, conceptualization is organized around what things do or is done to them, so functional information is crucial for both. At the same time, visual information, in the sense of the details of appearance, is less relevant for conceptualizing both living things and artifacts. ... Functional information, in the sense of what things do, is at the root of concept formation. (Mandler, 2004, p. 210)

In other words, there may be considerable scope for pursuing the claim that linguistic meaning highlights relational characteristics even for items that, at first sight, may seem ill suited.

By abstracting away from specific attributes, concerning which language users may have highly variable knowledge, the proposal also provides a basis for explaining how we can have a shared understanding of the words we use. We may differ widely in our ability to identify carnivores and draw specific inferences about them, but we can share in the more abstract relational notion of an animal that feeds on flesh. This is important because a significant characteristic of our language use, one

that is only rarely explicitly noted, is our apparently justified presumption that in general we share an understanding of the words we use.

Here I'm picking up on some ideas from Burge, who speaks of language users as having 'certain automatic reliabilities in understanding expressions' (2013, p. 347). Take almost any sentence from almost any communication, and while there may be considerable debate about the accuracy or clarity of what is expressed, and about the inferences that we are entitled to draw and the precise nuances of meaning intended by the author, it is also unusual to raise a query about the meaning of the individual words that are used in the utterance. We presuppose, and we seem to be warranted in doing so, that we share a basic understanding of the words that make up an utterance.

Of course, there are exceptions. Words such as *lepton* and *quark* may induce puzzlement. Queries can also arise in other special circumstances, such as where relatively unusual words are used. We might say 'I don't think that "lacuna" means what you think it does' if the speaker uses but appears to misunderstand the word. This type of query appears though to be rare. We may also query what may be called the interpretation of a phrase, as when we need further information to discern the intended referent of a phrase ('what bridge are you talking about?'), or when we query the applicability of a statement ('What do you mean it's not raining', and so on). With these examples we are not querying our basic shared understanding of the individual words used (bridge, rain); we are querying aspects of their application in a given situation.

It is perhaps quite remarkable that we appear to be entirely justified in presuming that we share a basic understanding of most of the words we use, however much we may debate aspects of interpretation.²¹ If so, an account of word meaning needs to be such that it supports the possibility of this shared understanding. We need some middle ground between a word taken purely as a pointer (which is not enough to support an understanding) and a full body of conceptual knowledge (which in general will not be shared between people). If it is correct that language use gives rise to a tendency to focus on shared relational structures, this provides a middle ground that would help to explain shared understanding.

The approach that I am arguing for has clear consequences for research into concepts and word meaning. The proposal indicates a criterion, not based on typicality or centrality or recognitional ability, for distinguishing between conceptual input that arises directly from knowledge of a word and conceptual input that is drawn from general knowledge (to which a word may though act as a pointer). The proposal also implies a constraint on the lexically driven aspect of compositional interpretation, in contrast to the more informationally rich extra-linguistic resources that may also influence our interpretations of utterances (see the discussion in Hogeweg & Vicente, 2020). Contextualist work on meaning has typically highlighted the role of contextually driven enrichments to interpretation, and one important idea that is sometimes considered is that word meanings may be underspecified in the sense of being in the 'wrong format' for directly constituting a part of a thought as expressed in an utterance (see Recanati, 2004). Such views require that inferential work be under-

²¹ I follow Burge (2013) in distinguishing an initial reliable 'comprehension' (Burge's term; I speak of basic understanding) from 'interpretation', which involves further inferential moves.

taken on the immediately encoded linguistic meaning, prior to a part of thought being expressed. The present proposal can be seen as providing an alternative to this theoretical tendency. It offers an account of linguistic meaning in which the immediately encoded material is already tailored to conveying an understanding, so can immediately appear as a constituent of that understanding without further inferential work.

Acknowledgements I would like to thank the two reviewers for this journal, whose comments and contributions led to a clear improvement in the presentation of the argument. I would also like to thank Robyn Carston, Agustin Vicente, Pauline Armary, and Antonin Thuns, for helpful and stimulating discussion.

Compliance with Ethical Standards There are no conflicts of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Acquaviva, P. (2014). Roots, concepts, and word structure: on the atoms of lexical semantics. In F. Rainer, F. Gardani, H. C. Luschützky, & W. U. Dressler (Eds.), *Morphology and Meaning* (pp. 49–70). John Benjamins
- Acquaviva, P. (2019). Categorization as noun construction: gender, number, and entity types. In E. Mathieu, M. Dali, & G. Zareikar (Eds.), *Gender and Noun Classification* (pp. 41–63). Oxford University Press
- Asmuth, J. A., & Gentner, D. (2005). Context sensitivity of relational nouns. *Proceedings of the Twenty-seventh Annual Meeting of the Cognitive Science Society*, 163–168
- Asmuth, J. A., & Gentner, D. (2017). Relational categories are more mutable than entity categories. *The Quarterly Journal of Experimental Psychology*, 70(10), 2007–2025
- Baumberger, C., Beisbart, C., & Brun, G. (2016). What is understanding? An overview of recent debates in epistemology and philosophy of science. In S. Grimm, C. Baumberger, & S. Ammon (Eds.), *Explaining Understanding. New perspectives from Epistemology and Philosophy of Science* (pp. 1–34). New York: Routledge
- Bierwisch, M., & Schreuder, R. (1992). From concepts to lexical items. *Cognition*, 42, 23–60
- Booij, G., & Lieber, R. (2004). On the paradigmatic nature of affixal semantics in English and Dutch. *Linguistics*, 42(2), 327–357
- Bourget, D. (2017). The role of consciousness in grasping and understanding. *Philosophy and Phenomenological Research*, 95(2), 285–318
- Brocher, A., Koenig, J. P., Mauner, G., & Foraker, S. (2018). About sharing and commitment: the retrieval of biased and balanced irregular polysemes. *Language Cognition and Neuroscience*, 33(4), 443–466
- Burge, T. (2013). Comprehension and interpretation. In T. Burge (Ed.), *Cognition through Understanding* (pp. 342–361). Oxford University Press
- Cooper, N. (2000). Understanding people. *Philosophy*, 75(3), 383–400
- Carston, R. (2016). The heterogeneity of procedural meaning. *Lingua*, 175–176
- Carston, R. (2019). Ad hoc concepts, polysemy and the lexicon. In K. Scott, B. Clark, & R. Carston (Eds.), *Relevance: Pragmatics and Interpretation* (pp. 150–162). Cambridge: Cambridge University Press
- Christie, S., & Gentner, D. (2014). Language helps children succeed on a classic analogy task. *Cognitive Science*, 38, 383–397

- Eliasmith, C. (2013). *How to Build a Brain: A Neural Architecture for Biological Cognition*. Oxford: Oxford University Press
- Embick, D. (2015). *The Morpheme: A Theoretical Introduction*. De Gruyter Mouton
- Fabregas, A., & Scalise, S. (2012). *Morphology: From Data to Theories*. Edinburgh University Press
- Fodor, J. (1998). *Concepts: Where Cognitive Science Went Wrong*. Oxford: Clarendon Press
- Fodor, J. (2008). *LOT2: The Language of Thought Revisited*. Oxford: Clarendon Press
- Franklin, R. (1981). Knowledge, belief and understanding. *The Philosophical Quarterly*, 31(124), 193–208
- Fricker, E. (1982). Semantic structure and speakers' understanding. *Proceedings of the Aristotelian Society*, 83, 49–66
- Gagné, C., & Spalding, T. (2007). The availability of noun properties during the interpretation of novel noun phrases. *Mental Lexicon*, 2, 241–260
- Gagné, C., & Spalding, T. (2013). Conceptual combination: the role of relational competition in the comprehension of modifier-noun phrases and noun-noun compounds. *Psychology of Learning and Motivation*, 59, 97–130
- Gagné, C., & Spalding, T. (2014). Subcategorization, not uncertainty, drives the modification effect. *Language Cognition and Neuroscience*, 29(10), 1283–1294
- Gentner, D. (2016). Language as cognitive tool kit: how language supports relational thought. *American Psychologist*, 71(8), 650–657
- Gentner, D., Anggoro, F. K., & Klibanoff, R. S. (2011). Structure-mapping and relational language support children's learning of relational categories. *Child Development*, 82(4), 1173–1188
- Gentner, D., & Hoyos, C. (2017). Analogy and abstraction. *Topics in Cognitive Science*, 1–22
- Gentner, D., & Smith, L. (2013). Analogical learning and reasoning. In D. Reisberg (Ed.), *The Oxford Handbook of Cognitive Psychology* (pp. 668–681). Oxford: Oxford University Press
- Glanzberg, M. (2011). Meaning, concepts, and the lexicon. *Croatian Journal of Philosophy*, 11(31), 3–31
- Glanzberg, M. (2014). Explanation and partiality in semantic theory. In A. Burgess, & B. Sherman (Eds.), *Metasemantics: New Essays on the Foundations of Meaning* (pp. 259–292). Oxford: Oxford University Press
- Glanzberg, M. (2018). Lexical meaning, concepts, and the metasemantics of predicates. In D. Ball, & B. Rabern (Eds.), *The Science of Meaning: Essays on the Metatheory of Natural Language Semantics* (pp. 97–225). Oxford University Press
- Goldwater, M. B., & Markman, A. B. (2011). Categorizing entities by common role. *Psychonomic Bulletin and Review*, 18, 406–413
- Hampton, J. A. (2015). Categories, prototypes and exemplars. In N. Riemer (Ed.), *The Routledge Handbook of Semantics* (pp. 125–141). London: Routledge
- Hampton, J. A., Passanisi, A., & Jönsson, M. L. (2011). The modifier effect and property mutability. *Journal of Memory and Language*, 64(3), 233–248
- Haiman, J. (1980). Dictionaries and encyclopedias. *Lingua*, 50, 329–357
- Hogeweg, L., & Vicente, A. (2020). On the nature of the lexicon: the status of rich lexical meanings. *Journal of Linguistics*, 865–891
- Katamba, F., & Stonham, J. (1993/2006). *Morphology* (2nd edition). Palgrave Macmillan
- Keil, F. C. (1989). *Concepts, Kinds, and Cognitive Development*. Cambridge, MA: MIT Press
- Kvanvig, J. (2003). *The Value of Knowledge and the Pursuit of Understanding*. New York: Cambridge University Press
- Levin, B., & Rappaport Hovav, M. (2005). *Argument Realization*. Cambridge: Cambridge University Press
- Löhr, G. (2020). Concepts and categorization: do philosophers and psychologists theorize about different things? *Synthese*, 197, 2717–2791
- Longworth, G. (2018). Understanding what was said. *Synthese*, 195, 815–834
- Lupyan, G. (2012). What do words do? Toward a theory of language-augmented thought. In B. H. Ross (Ed.), *The Psychology of Learning and Motivation* (57 vol., pp. 256–297). Waltham, MA: Academic Press
- Lupyan, G., & Lewis, M. (2019). From words as mappings to words-as-cues: the role of language in semantic knowledge. *Language Cognition and Neuroscience*, 34(10), 1319–1337
- Machery, E. (2009). *Doing Without Concepts*. Oxford: Oxford University Press
- Mandler, J. (2004). *The Foundations of Mind: Origins of Conceptual Thought*. Oxford: Oxford University Press
- Millikan, R. (2000). *On Clear and Confused Ideas*. Cambridge: Cambridge University Press
- Namy, L., & Gentner, D. (2002). Making a silk purse out of sow's ears: young children's use of comparison in category learning. *Journal of Experimental Psychology: General*, 131(1), 5–15

- Olsen, S. (2012). Semantics of compounds. In C. Maienborn, von K. Heusinger, & P. Portner (Eds.), *Semantics: An International Handbook of Natural Language Meaning* (3 vol., pp. 2120–2150). De Gruyter Mouton
- Pettit, D. (2002). Why knowledge is unnecessary for understanding language. *Mind*, 111(443), 519–550
- Perry, L., & Lupyan, G. (2017). Recognizing a zebra from its stripes and the stripes from “zebra”: the role of verbal labels in selecting category relevant information. *Language Cognition and Neuroscience*, 32(8), 925–943
- Peeters, B. (2000). Setting the scene. Recent milestones in the lexicon-encyclopaedia debate. In B. Peeters (Ed.), *Current Research in the Semantics/Pragmatics Interface* (vol. 5, pp. 1–52). Oxford: Elsevier Science
- Pietroski, P. (2018). *Conjoining Meanings: Semantics Without Truth Values*. Oxford: Oxford University Press
- Pritchard, T. (2017). Knowing the meaning of a word: shared psychological states and the determination of extensions. *Mind and Language*, 32(1), 101–121
- Pritchard, T. (2019). Analogical cognition: an insight into word meaning. *Review of Philosophy and Psychology*, 10, 587–607
- Putnam, H. (1975). The meaning of ‘meaning’. In K. Gunderson (Ed.), *Language, Mind, and Knowledge* (pp. 131–193). Minneapolis: University of Minnesota Press
- Quilty-Dunn, J. (2021). Polysemy and thought: toward a generative theory of concepts. *Mind and Language*, 36, 158–185
- Ramchand, G. (2008). *Verb Meaning and the Lexicon: A First Phase Syntax*. Cambridge University Press
- Rappaport Hovav, M., & Levin, B. (1992). -ER nominals: implications for the theory of argument structure. In T. Stowell, & E. Wherli (Eds.), *Syntax and Semantics, Vol. XXVI: Syntax and the Lexicon* (pp. 127–153). New York: Academic Press
- Recanati, F. (2004). *Literal Meaning*. Cambridge University Press
- Recanati, F. (2017). Contextualism and polysemy. *Dialectica*, 71(3), 379–397
- Riggs, W. (2003). Understanding ‘Virtue’ and the Virtue of Understanding. In M. DePaul, & L. Zagzebski (Eds.), *Intellectual Virtue* (pp. 203–226). Oxford: Clarendon Press
- Ruhl, C. (1989). *On Monosemy: A Study in Linguistic Semantics*. New York: State University of New York Press
- Vicente, A. (2018). Polysemy and word meaning: an account of lexical meaning for different kinds of content words. *Philosophical Studies*, 175, 947–968
- Zagzebski, L. (2001). Recovering understanding. In M. Steup (Ed.), *Knowledge, Truth and Duty* (pp. 235–251). New York: Oxford University Press

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.