

Analogical cognition and understanding a word*

Timothy Pritchard

Abstract

Gentner describes a notion of ‘analogical cognition’ that could play a significant role in elucidating what is involved in understanding a word. Gentner’s work has not, though, had much or any impact in linguistics or the philosophy of language. I explain key features of Gentner’s notion, and I argue that it explains how word understanding can be stable, specific, and shared, and how words can contribute to cognition as opposed to just being a means of conveying thoughts.

Keywords: analogical cognition, Gentner, word meaning, understanding, good enough interpretation

1 Introduction

Speakers of a language appear to have a good understanding of most of the words they use. They have an easy and pre-reflective familiarity with individual words, and they draw on this both when speaking and when listening. Even if the thoughts that speakers use combinations of words to convey can be difficult to understand, this difficulty is not typically felt at the level of the individual words themselves. I take these appearances to be well founded and to legitimate the following question: can we find a way to characterize how we understand individual words? The intended focus of the question is not on what we might use a word or words to convey in a given instance, but on our prior understanding of the words, the understanding that we draw on when we communicate. There appears to be little confidence among theorists about how to approach this question; for example, while Recanati (2004) does discuss various options, there is little by way of development of, or debate between, these options.

In this paper, I offer a specific proposal that is based on Gentner’s work on analogical cognition. Her work provides, I believe, resources for making a significant contribution to the debate. T. Bach (2012) has commented that the study of analogical cognition has had little impact on the philosophy of mind and language, and I agree with his opinion that this could be an ‘expensive oversight’. Insight from Gentner’s work helps to motivate and justify the claim that our understanding of particular words is stable, shared, and quite specific. It enables us to delimit a specifically linguistic level of understanding as opposed to more general types of understanding, and it helps to explain how the use of words makes a special contribution to cognition. The account stands critically over against approaches that, inspired by reflection on context sensitivity, have tended to offer indirect, open-ended, and non-specific accounts of the contribution that words bring to utterances.

My preferred terminology is to phrase the debate in terms of understanding rather than in terms of (word) meaning. Notions of meaning often relate to the question of how to represent what a given use of language has conveyed, and this is not always helpful for considering our grasp of individual words. I would prefer to take our understanding of individual words as prior

* I would in particular like to thank Andrew Tolmie for indirectly pointing me towards Gentner’s work. Also, for helpful discussion and corrections, Josie Bowerman, Robyn Carston, Léna Kervran, Nathan Klinedinst, Lewis Pollock.

in order of explanation to accounts of meaning, but nothing turns on this distinction in this paper. Sometimes, though, it will be more natural to speak of word meaning, in particular when this reflects the terminology of other theorists. We might think of a word meaning as the corollary to what is understood in the understanding of a word, though it is useful to note that this takes us away from a more natural emphasis in which ‘meaning’ relates to what is conveyed in the use of language.

So, while understanding is my preferred term, I will also at times speak of word meanings. Either way, the focus is on words considered in a sense prior to their use in utterances. I ignore aspects of understanding that relate specifically to syntactic patterns and morphology, so the focus is more towards what might be described as an underlying or root understanding that may (but need not) recur across different word forms. For example, the same root understanding typically recurs in ‘capitulate’ and ‘capitulation’ but not typically in ‘capitulate’ and ‘recapitulate’.¹

In section 2, I describe two characteristics of language that an account of word understanding should help to explain. In sections 3 and 4, I describe the relevant aspects of Gentner’s notion of analogical cognition, and in section 5, I propose an account of word understanding that is based on this. In section 6, I discuss how the account fits the two characteristics of language described in section 2. In section 7, I consider how the contribution words make to cognition challenges some accounts of the basis for word use.

2 Two characteristics of language

Two characteristics of language that an account of word understanding should help to explain are context sensitivity and the phenomenological immediacy of our experience of meaning when we hear language. There are other important characteristics, some of which I shall mention (such as the ease with which children learn words), others of which I do not directly address (such as compositionality). The two I focus on are particularly pertinent for the account I describe.

Emphasis on a stable, shared, and quite specific understanding (or, stable and specific word meaning) is often seen as problematic. Reflection on language use has suggested to many theorists that we need to treat words’ contributions to utterances in an open-ended and indirect fashion, and that it may even be wrong to think of words as having any specifiable meaning attached to them at all (for overview, see Carston (2012)). These suggestions have arisen from consideration of the flexibility with which words are used, for example as reflected in polysemy or in nuances of interpretation that are sometimes described as involving a modulation of word meanings (Wilson & Carston, 2007). The contribution that words make has been described with terms such as ‘constraint’ (Travis, 2000), ‘clue’ or ‘pointer to conceptual space’ (Carston, 2002), something that points to ‘indefinitely many notions or concepts’ (Sperber and Wilson, 1998), ‘instruction for creating a concept from available mental resources’ (Pietroski, 2005), ‘polysemy complexes’ (Carston, 2016).²

¹ H. Borer (2005) gives an important place to *word* roots (after the pattern of tri-consonantal Hebrew roots), a level of word analysis that is prior to any particular morphological implementation. She does not though speak of a (differently grained) notion of root understanding.

² Borg, who dissents from this general approach, characterizes the view as one that treats word contributions ‘as open-ended, web-like things which stand in need of contextual precisification prior to fixing their input to larger linguistic units’ (2012, xvii). Her own view is that word contributions are ‘discrete, probably atomistic, blob-like things’. The account I propose supports a move away from the open-ended to something more specific.

These suggestions indicate roles, and there has been discussion of what plays these roles. Carston has considered (though now rejects) the view that word meanings could be ‘schemas’ (2002; 2012; 2013). A schema is envisaged as a highly abstract representation that captures merely an outline of the different meanings that a word can be used to express. The schema is, as it were, the skeletal starting point for interpretation, with a fleshing out needed before we attain to an interpretation that will be typical of a given use of a word. Alternatively, word meanings have been characterized as containing more than is required for interpretation. Rather than needing to be fleshed out, a filtering process is required, in which the relevant interpretation is selected (see Recanati (2004), p.140).

More radical suggestions have been proposed, in which word meaning is eschewed altogether. Drawing on Wittgenstein, Recanati (2004) describes (without fully endorsing) a ‘meaning eliminativist’ view. On this view, word use is explained via knowledge of previous uses of a word (source situations), with further use of the word licensed when a novel situation is appropriately similar to those previous uses. Rayo (2014) likewise proposes that rather than thinking in terms of linguistic meanings, it is better to see speakers as associating words with a ‘grab bag’ of mental items: memories, mental images, pieces of encyclopaedic information, and so on. What is conveyed by a word arises from some process that is based on language users having access to this grab bag of items.³

The suggestions reflect the claim that a word is but a pointer, or clue, towards the relevant interpretation of a use of a word. There may or may not be a specifiable meaning to a word, but either way the word acts as an initial input that requires further work before the required interpretation can be recovered. Language users’ grasp of a word does not directly correspond with what a word is used to express on an occasion of use.

There is, though, besides context sensitivity, another and apparently contrasting aspect to our use of language. This is the observation that, on hearing words, understanding arises in such a direct and involuntary manner that it is almost perceptual in character. Fricker (2003) describes how when we hear language, we seem to experience the meaning of the words in the very hearing of the utterance itself; the utterance is experienced as a ‘semantically laden event’. Fricker calls this an understanding *experience*, a quasi-perception. It seems as if the meaning is an intrinsic property of the utterance.

[W]hen one understands an utterance of a sentence, the immediate object of one’s aural experience is not a mere-burst of sound, nor merely a syntactic entity. The phenomenal given in normal language perception is an utterance perceived as syntactically structured and semantically loaded. One hears the words and one hears what they mean—one hears them as expressing a certain proposition. (2003, pp. 337–8)

While there is debate as to whether the parallel with perception is justified, the broad description of the phenomenology is not controversial: at the conscious personal level, meanings seem, at least very often, to directly accompany the words that are used. Further, we cannot choose to avoid uptake of meaning when we hear language. As Pettit (2010) describes it:

Our awareness of the content of speech is typically immediate and unreflective. Indeed, the content of speech in your native language can stubbornly intrude upon

³ Arguably, meaning eliminativist views are merely more thorough-going presentations of the ‘rich’ view. Both are characterized by a need to filter out the relevant use from an array of interpretive material (see Carston (2012), p. 620, on Cohen).

your mental life when you would rather it not, as for example when someone is publicly sharing personal details while talking on the phone. You can plug your ears or try to attend to something else, but you cannot simply block out the content of the speech.

Fricker and Pettit are careful to note that these personal-level observations about the immediacy of meaning may mask complex underlying processes. But it is also reasonable to require that an account of word understanding should help to explain the phenomenon. Prima facie, it is not clear that talking of words as pointers or clues or constraints is best suited for providing an explanation, as they are approaches to words that stress indirectness and inference rather than involuntary directness.

A parallel point can be made with respect to the speed with which words are processed. Keil and Kominsky (2015) express this (in terms of ‘concepts’):

concepts, as manifested through words, must be processed at rates compatible with both producing and understanding speech and the reading of text. When a reader seems to fully comprehend a text at three hundred words a minute, how are hundreds of concepts accessed, deployed, and interrelated each minute? ... [F]or concepts to work in this way, they must impose a very low cognitive load and not require much inferential machinery in real time. It would not be possible to deploy concepts so quickly and effortlessly if they required extensive processing of an elaborate web of beliefs.

If words are pointers or clues, this may imply a cognitive load that is hard to reconcile with the speed of processing. The immediacy and speed of our uptake of meaning may more naturally suggest that individual words have specific and easily understood meanings, which directly, and not via inference, inform our understanding of the utterances in which those words are used. This of course raises the problem of how to explain the evidence that word use is highly flexible and context sensitive.

3 Analogical cognition⁴

A striking feature of human cognition is our ability to handle higher-order relations. Consider the following matching task, which requires a match to be made between the appearance of the base sample (the letter A) and one of the target letters:

Base	Target
A	A B

An ability to evaluate the perceptual similarity between two stimuli has been described as ‘clearly the sine qua non of biological cognition’ (Penn, Holyoak, & Povinelli, 2008), and the ability to match the Base A to the Target sample A is not limited to humans. The matching is based on comparing features of the various tokens. But consider the following different type of matching task:

⁴ I focus on those aspects of Gentner’s work that are particularly relevant to the understanding of individual words. Her work on analogy goes beyond this, including consideration of various learning and reasoning abilities.

Base	Target
AA	BB CD
XYX	ZZW UVU

This is a relational-match-to-sample test, and it is not solved by looking for perceptual similarity between the individual letters in the Base and the Target. Instead, the matching task looks first to a relation that holds between the individual tokens that make up each sample, with matching based on recurrence of that relation. AA matches to BB, and not to CD, because the same relation holds between A and A (namely, sameness) as holds between B and B. Similarly, XYX matches to UVU, not to ZZW, because both display the relation of symmetry.

The examples illustrate the idea of a structure into which individual items enter, structure that, at least for humans, invites a form of cognition that is based on an appreciation of that structure. Even if the items themselves differ considerably, they can still enter into the same relations.

[T]he corresponding objects in the base and target need not resemble each other; what is important is that they hold like roles in the matching systems of relations.

(Gentner, 2010, p. 752)

A railway engine can have the same relation to a carriage as a tractor does to a cart. Humans are sensitive to this shared structure and express it with terms such as ‘pulls’. A concrete block that prevents a car from progressing down a road manifests a situation that shares structure with a book that prevents a marble from rolling further on a table. Humans express this with terms such as ‘obstacle’.

Analogical ability, also called relational cognition, is the ability to perceive like relational structures across different contexts. There is debate as to whether analogical cognition is limited to humans (see Penn et al. (2008)), but there is general agreement that human thought is specially characterized by an ability to engage in this kind of relational matching and that this ability greatly exceeds any relational matching ability that some animals might turn out to possess. It is ubiquitous in human cognition and has been described as the hallmark of human intellectual ability:

Analogical ability ... is a core mechanism of human cognition. The ability to perceive and use purely relational similarity is a major contributor—arguably *the* major contributor—to our species’ remarkable mental powers.

(Gentner & Smith, 2013)

Gentner and her colleagues have shown that this analogical ability is in particular helped by the process of *comparing* different situations or entities (Gentner 2005; 2010). Research shows that when we compare different items, we preferentially highlight not the perceptual similarities but the relational commonalities. According to what Gentner calls structure-mapping theory, this is because when people compare two items or situations, they implicitly seek to find structural correspondences (Gentner & Markman, 1997; Gentner, 2003). This directs attention away from consideration of whether the compared items share an overall (and in particular perceptual) likeness, and towards consideration of shared structure. This ‘relational shift’ gives rise to a form of learning (schema abstraction) in which the common structure is extracted.

This relational shift has been illustrated in young children. In a classic search task,

children watched an experimenter hide a toy in a small model room and had to find another toy hidden ‘in the same place’ in the full-sized room. Children under 3 years of age did very poorly at this task. Building on these results, [a further study] showed (a) if 2.5 year olds are given a similar-scale search task (so that model and room are highly similar) they perform very well (67% correct searches); and (b) when these children are brought back the next day and tested in the standard model-room task (small model and full-sized room), they are far more successful than age-mates who did not receive the highly similar pair first (35% correct). (Gentner & Smith, 2013)

When the model and the room were similar in size, this facilitated the task of looking for an item ‘in the same place’ in both rooms. This also led, by the implicit comparison undertaken of the two rooms, to an appreciation that both rooms shared a structure and that sameness of place is evaluated with respect to that structure (and not, for example, with respect to local perceptual similarities). This is shown in the children displaying improved ability in evaluating sameness of place when confronted with the different sized rooms. Importantly, learning from highly similar things did not lead the children to put an increased emphasis on perceptual similarity as the criterion for sameness of place. Instead, the effect of similarity between the rooms was to make the act of comparing easier. Gentner calls this ‘progressive alignment’, with the comparison of similar things acting as ‘training wheels’ (Gentner, 2010) that help the child to learn of relational matches that can subsequently be used to find samenesses across perceptually dissimilar items.

In another study, children (aged 3.5 years) watched the experimenter hide a toy in a small model room (the Hiding room), and then tried to find the toy hidden ‘in the same place’ in a second model room (the Finding room). The two rooms contained the same types of furniture (bed, table, etc.) in the same configuration, but the specific shapes of the furniture in the Finding room differed from the shapes of the furniture in the Hiding room.

Before engaging in the task, all the children were shown the Hiding room along with another highly similar room (identical except for color). Half the children saw the two rooms together and were encouraged to compare them; the other half discussed each room separately. Children who had compared the rooms were significantly more likely to correctly locate the toy in the Finding room than those who had experienced the rooms separately. Thus, comparing two nearly identical rooms facilitated children’s ability to map their common spatial relational structure to a relationally similar but surface-dissimilar target.

(Gentner, 2010, pp. 759–60)

That is to say, when the children were encouraged to compare the two highly similar rooms, they showed a tendency to pick up on a relational specification of the location of the object (viz., looking to the same relative position) and to be less influenced by the changing shapes of the items of furniture involved.

For adults, likewise, eliciting comparisons helps induce a focus on abstract relational commonalities. Business school students who compared two negotiation scenarios were over twice as likely to transfer the negotiation strategy to an analogous test negotiation as were those who studied the same two scenarios separately (Gentner, 2010, p. 758).

In summary, these and many other studies show that comparing items or situations leads human thinkers to emphasize structural commonalities between those items or situations. Comparison renders the common structure more salient, and leads to a form of learning in which the common structure is extracted. This underpins analogical cognition, which is the

ability to perceive the same structural relations across different situations. Gentner has also shown that analogical cognition is closely related to our use of words.

4 Analogy and language

When discussing what it is that makes humans so smart, Gentner does not just talk about relational cognition; she combines it with our use of language. What makes us smart is ‘a mutual bootstrapping system composed of (a) our exceptional capacity for relational cognition and (b) symbolic systems that augment this capacity’ (Gentner, 2010). Language supports our analogical reasoning abilities.

Experimental work provides plenty of evidence for the important role that a linguistic label itself can play in aiding a shift to relational thinking. For example, Gentner (2003) reports an experiment in which children (age 3 to 5) were shown two groups of three objects (the child’s set and the experimenter’s set) arranged according to size. For example, in a simple case there might be three differently sized cups. Children watched as the experimenter hid a sticker under an object in the experimenter’s triad. The children were told they could find their sticker by looking ‘in the same place’ in their own group of objects. The criterion for the same place was relative size; so if, for example, the experimenter placed the sticker under the largest object, the child should choose the largest object in the child’s group of three objects. The children were indirectly taught about this criterion in 14 trial attempts. In each trial, the children would suggest an answer and then be shown the correct response.

The children were particularly unsuccessful in the task when the two groups of objects contained the same items but differently relationally placed. For example, suppose that the largest item (under which the sticker is placed) in the experimenter’s triad was a cup, and the children’s group of three objects had a cup but it occurred in the middle or smallest position. With this type of condition, instead of choosing the largest item in their own group of objects, children would often choose the cup, answering on the basis of perceptual similarity. In other words, in spite of repeated trials in which the relational criterion was used, the children were, under certain conditions, quite easily side-tracked away from a focus on this criterion and towards perceptual similarity.

This dramatically changed when the objects in the groups received linguistic labels that related to the relational dimension. For example, the labels could be ‘big’, ‘little’, ‘tiny’. With this change to the conditions of the experiment, there was a massive shift to relational thinking and away from the use of perceptual similarity as a guide. The effect persisted over time, with the label-trained children performing well on future tests even when the labels were no longer used. In other words, use of the linguistic labels helped to shift evaluation of what it is to be ‘in the same place’. Gentner suggests that children learn to become analogical thinkers, with language facilitating this relational shift.

Similar results have been found for spatial terms such as ‘top’, ‘middle’, ‘bottom’ (Loewenstein & Gentner, 2004). When the position of an item is described as on ‘top’ of a box (for example), children very much more readily recognized that the relational position was the relevant criterion when asked to locate a test item that is in the ‘same’ position with respect to a different box. When the position of the item was described simply as ‘here’, children were less ready to recognize or act on the relational criterion.

Christie and Gentner (2014) report an experiment in which children (average age 2.5 years) were tested on a relational-match-to-sample test. Initially, even after feedback, the children did not pick up on the relational criterion (answering at only chance levels). But when the children were given a novel linguistic label for the standard (a ‘truffett’), and told to apply that label to the choice of alternatives, there was a shift to the relational criterion. This effect is

described as ‘quite remarkable’ given that the word was novel and the children in this condition received no training with the novel word or feedback. It appears that when a matching task is presented in terms of applying a term, this facilitates access to a relational criterion amongst very young children for whom that criterion is otherwise not particularly accessible.

Similarly, Gentner (2005) reports an experiment that used the novel label ‘blick’. In the first condition, children were told that ‘The knife is the blick for the watermelon. The axe is the blick for the tree’. They were then asked, ‘What would be the blick for the paper?’, and were offered three options: a thematic response (pencil), a perceptually similar response (another piece of paper), and a relational response (scissors). In a second condition, instead of using this novel ‘blick’ label, the children were told that the knife goes with the watermelon in ‘the same way’ that the axe goes with the tree. They were then asked what would go with the paper in ‘the same way’. For 4 and 6 year olds, the relational response was more likely in the blick condition than in this second condition.

An inviting explanation of this is that linguistic terms, being applicable to multiple different particular situations, are naturally aligned with the relational commonalities that comparison of situations encourages. Labelling a situation will thereby prompt a search for some relevant structural relation (Gentner, 2005). This gives rise to a process of abstraction, the use of a label helping to ‘de-situate’ the relational structure from the initial context (Gentner, 2003). Understanding of this structure underpins subsequent use of the term, usage which is not based on expectations of perceptual similarity.

Use of a linguistic label also helps preserve the abstraction in memory and renders it more accessible for future use (Gentner calls this ‘reification’). “The most obvious instance of this increased accessibility is that a named pattern can be re-invoked by using the name” (Gentner & Christie, 2010). This may seem a trivial point, but it is highly significant. As I will briefly discuss in section 7, there is evidence that without the use of linguistic labels, these relational patterns are not as accessible to cognition as more obvious perceptual similarities. Language is, *par excellence*, a device in which labels are re-used across multiple items and situations that may show little or nothing by way of perceptual similarity. Rather than simply being a medium through which analogical cognition is recorded, the suggestion here is that language itself plays a constitutive role in facilitating this type of cognition.

5 A hypothesis for word understanding

The evidence supports the following hypothesis: at least for many words, for a word to be understood is for the word to trigger a focus on a structural relation. This focus underpins the matching capacity that is pervasively illustrated in our use of a word. The sheer complexity of word use will always challenge any simply drawn suggestion. But the hypothesis is well motivated and has various interesting characteristics, which I discuss in the rest of the paper.

The account dissociates the basis for word use from whether or not items/situations share properties, in particular perceptual likenesses (‘surface’ similarities, as Gentner sometimes puts it). This is massively confirmed by how we use words: we effortlessly, and without thinking of the dissimilarities, apply a single term across situations that can differ widely in the particular features that make up those items or situations. We do not think of the differences when we describe an elastic band, an electric device, an internal combustion engine, part of a biological organism, or even a person, as a ‘motor’. Use of the single word ‘motor’ focuses our attention

on a shared structural relation (roughly, something that imparts motion).⁵ We can recognize when a negotiation is occurring, whether that is between a mother and her child, a buyer and seller, between the members of a committee, or between nations. These situations differ in many ways, but they are the same in respect of manifesting a given structure (roughly, a discussion aimed at reaching an agreement). The challenge is not to think up such examples, but to think up examples where word use is based on something like perceptual similarity instead.

Gentner and Kurtz (2005) observe that a relational status is plausibly attributed to a great number of the nouns that we use. They contrast relational categories, whose members ‘cohere on the basis of a core relationship’, with what they call entity categories, whose members ‘share many intrinsic properties’ (which will often be perceptual, see Gentner (2005)).⁶ Words such as ‘tulip’ and ‘camel’, they suggest, name entity categories. Gentner and Kurtz do not, though, distinguish the weaker claim that some words are principally used of items that happen to display overall likeness from the stronger and more relevant claim that the understanding of some words is constituted in terms of a grasp of a likeness that must hold between members of the category. Just because members of a category happen to be perceptually similar, for example, it does not follow that this similarity plays a constitutive role in our understanding of the word for that category. For example, books used to be far more uniform in appearance and methods of construction than they are today; but modern production techniques, and in particular the shift to electronic books, show that these uniformities were not actually determinative for how the word ‘book’ was understood. Further, some of Gentner and Kurtz’s proposed examples of terms for entity categories, such as ‘household appliance’ and ‘musical instrument’, are clearly inadequate: neither are used to name items that need share overall likeness, and both have plausible structural explanations (pianos and triangles both play a role in enabling music to be produced; kettles and irons are both manufactured products used for particular household operations). The important point, though, independently of what we think of the proposed entity categories, is that for at least very many words in regular use a structural explanation is plausible.

This claim is consistent with the observation that, in practice, surface similarities can be used to help guide our use of words (in particular for categories whose members do display general overall likeness). Many studies have shown the importance of shape as a cue for choosing the appropriate item to label with a given noun.⁷ This is adequately explained, on the current hypothesis, as due to a correlation that often holds between shape and the relations into which items enter. This is particularly clear with the case of function. If it looks like a hammer or a bicycle, an item can probably play the roles of a hammer or bicycle. But this does not show that the understanding of the label is constituted in terms of the shape of an item. Abilities to sort, and make inferences concerning, a given category go beyond the input from word understanding. The input from word understanding provides (by hypothesis) only a limited basis for sorting and inferential abilities; these restricted abilities will draw on consequences that follow from the specific structural feature that a word picks up on. While I may use shape to identify that an item is a motor, this ability does not arise from my understanding of the word

⁵ I take this rough description from the OED. Lexicographers have often provided a wealth of suggestions as to how to characterize structural relations.

⁶ Use of the qualifier ‘intrinsic’ is meant to indicate a contrast with ‘relational’, but it is an unfortunate term to use. It is intrinsic to a barrier to be an item that prevents something from moving. I take the contrast to be between the claim that members are alike in some overall fashion and the claim that membership is based on a structural relation. As I point out, these are not mutually exclusive.

⁷ On the shape-bias effect, see for example the special issue *Developmental Science*, 11(2), 2008.

‘motor’, which, on the current proposal, is limited to the relation of (roughly) something that imparts motion. Knowledge of shape, and other similar types of knowledge, belong to general knowledge about motors, and are not part of linguistic understanding.⁸

Use of the word ‘mouse’ for the computer accessory was initially motivated, according to its inventor Doug Englebart, on grounds of similarity in appearance to a mouse. Interestingly, this is a case of a word acquiring a new meaning (homonymy) rather than an extension based on what the word ‘mouse’ originally meant. Further, while shape was the motivating factor for using this label, shape is not the basis for further use of the computer word ‘mouse’. Computer mice often no longer have tails (the wire), but that is not important for the extension of the word, which is based on a relation that the item enters into between a user and a computer.

By providing a limited scope to word understanding, the account helps to explain why children can learn so many words with such apparent ease and rapidity. Word understanding is based on a quite specific appreciation of a particular structure, one that is available even to young children. By hypothesis, a child can have as good a linguistic understanding of the word ‘motor’ as an expert mechanic does.

Before considering further characteristics of the approach, I will finish this section with consideration of what is probably going to be regarded as a key difficulty for the account.

Wittgenstein (1953, §66) famously suggested that for at least some words, such as ‘Spiel’ (game, that which is played)⁹, there is no common factor that explains the use of the word across the range of instances to which it is applied. Wittgenstein’s example is indeed highly suggestive, and it is not an easy task to think up an account that accurately reflects or predicts the range of situations to which we apply the term and also the range of situations to which we do not apply the term. In English usage, ‘game’ applies to football and cricket, but not (it seems) to athletics events; it applies to Trivial Pursuits but perhaps not to pub quizzes; it may apply to some types of training exercise (War games) but not to others. A particularly useful feature of the example is that there is no compelling reason to appeal either to homonymy, or even to polysemy,¹⁰ when explaining at least the majority of our uses of ‘game’. So, according to the style of account that I am proposing, there ought to be some relatively simple relation that is characteristic of the instances to which we apply the term ‘game’.

While Wittgenstein’s example does show the difficulty of the task, he does not show that it is futile to look for such an account. Three points can be made.

First, he doesn’t actually show that ‘game’ has no unifying explanation for application. He merely lists a range of suggestions as to what the commonality (or commonalities) might be—such as being amusing, or involving winning and losing—and after a brief attempt concludes that no such commonality is to be found. Arguably, Wittgenstein was looking for some bundle of individual features that might hold of all and only games, rather than to a structural relation that is characteristic of games as such.¹¹

Second, part of Wittgenstein’s own positive proposal for how terms are applied (insofar as it is such a proposal) appeals to ‘family resemblances’. It is significant that this idea was

⁸ In agreement with Fodor (1998; 2008), the account indicates that abilities are not constitutive of concepts (in my terms, word understandings), though, unlike Fodor, some abilities follow directly from the limited sphere of word understanding. Fodor empties the lexicon of any such sphere of understanding, and reduces the connection between lexical concept possession and abilities to an association that holds by virtue of lexical concepts acting as pointers to mental files, in which all the epistemically relevant material is found.

⁹ See Hoyningen-Huens (2015). ‘Spiel’ has a wider range of application than the English word ‘game’.

¹⁰ Polysemy does not have a clear definition, but my point is that a vast number of uses of ‘game’ do not need to be qualified as showing variations of sense, even if the games it applies to vary much in characteristics.

¹¹ Hoyningen-Huens (2015) suggests that Wittgenstein wasn’t looking at the appropriate level of abstraction, and he thinks that an account of ‘Spiel’ can in fact be given.

illustrated by a largely perceptual example (the build and features of persons in a family, (1953, §67)), and the focus is on comparing some appropriate number of features of different items. As we have seen, there is compelling evidence that looking to structural relations is important for at least many words, and this is a focus that removes attention from making a comparison between the features of items.

Third, the context of Wittgenstein's discussion is that of showing that we cannot explain word use by appeal to a putative meaning that somehow determines, by itself, the items or situations of which the word can be truly used. This is an important philosophical point, but it is not one that the appeal to analogical cognition need be seen as inconsistent with. The focus with analogical cognition is on a commonality of understanding that can be specified at the linguistic level; the focus is not on attempting to delineate sets of defining marks that would have the role of dividing up items into fixed extensions. Usage indicates that a marble can be correctly described as a barrier, and correctly described as not a barrier. This is predicted by the analogical approach: whether the structural relation is realized will depend on what enters into the relation with the marble.¹²

Even if appeal to analogical cognition is broadly correct, there are reasons for thinking that it may often prove difficult to pin down, in explicit terms, our understanding of a word. In practice, word use is influenced by how neighbouring words are used. Language aids cognition not only by providing individual labels, but by providing a multitude of related labels, creating a complex multi-dependent framework of word usage that both induces a fineness of grain in our understanding of a given word and also leads to contingencies of usage that may be beyond the explanatory reach of theory. Use of terms such as 'athletics' (for track and field events), or 'swimming' (for pool events), seems to block the use of 'game' for these events, whereas there is no such block for football, cricket, and hockey. On the other hand, we also speak of the Olympic Games (a point noted by Hoyningen-Huens (2015)), which illustrates the contingencies that can be displayed in word use. In terms of structural relations between user and computer, a track pad on a laptop need hardly be distinguished from a mouse. But we do make a distinction, and in this case it seems plausible to suppose that appreciation of the difference in how the items are manipulated enters into our understanding of 'mouse' and 'track pad'.

Related words, therefore, may lead to usage of a given word being influenced by a potentially complex array of factors. This is consistent, though, with the claim that looking to structural relations provides an explanation for the primary basis on which word understanding is founded.

6 Understanding and context sensitivity

One response to the phenomenon of context sensitivity has been to suggest that word meanings are 'underspecified'. That is to say, what a word means does not itself fully determine the content that a use of the word may come to convey. Rather, the word meaning provides some sort of starting point, from which a particular interpretation may be constructed.

An approach to word understanding based on analogical cognition has an element of underspecification, and in this respect is well able to deal with at least some types of context sensitivity. Because the emphasis is on a structural relation, and not the specific features of the items that manifest this relation, nothing is determined in the account as to how a particular

¹² The point is anticipated in Plato's discussion of Forms. Forms can be understood in terms of structure (see Prior (1983)), and this helps explain the phenomenon of the 'compresence of opposites', in which a predicate both applies, and does not apply, to a particular item (see Pritchard (forthcoming)).

item need be instantiated in a given instance. Understanding the word ‘obstacle’ or ‘motor’ does not involve any determination for the type of obstacle or motor it is or for any of their features—beyond that they will help satisfy a given relational role. The way in which something is an obstacle, for example, or the perspective from which it is an obstacle, is not specified at all.

So, on hearing the utterance ‘The motor pulled the obstacle out of the path’, our linguistic understanding presents us with a claim in which an item that causes movement stands in the pulling relation to an item that plays a preventing role on a terrain that enables movement between points (a rough characterization of ‘path’).

But if so much is left unspecified, can we explain the phenomenology, in which meanings appear to be directly apprehended in the hearing of words? The answer, I believe, is yes, and relates both to the fact that the account does suggest a *specific* meaning (as opposed to some vague pointer to conceptual space) and on the observation that understanding is not an all or nothing affair—language users often appear to make do with merely ‘good enough’ understandings. We need to allow for the possibility that the kind of understanding that is immediately experienced on hearing words has no more than the specificity contained in the semantics of the words used themselves, even if further levels of representation are always possible.

Of particular interest in this respect is evidence that language users do not seek for fully specified interpretations of utterances, but make do with ‘good enough’ representations for the purposes at hand (see Ferreira, Bailey, and Ferraro (2002); Stanford and Surt (2002), from which the following examples are taken). For example, consider the pronouns in the following sentences:

- (1) Mary bought a brand new Hitachi radio.
- (2) It was in Selfridge’s window.
- (3) Later, when Joan saw it, she too decided it would be a good purchase.

The referent of *it* in (2) is not fully specified. Did Mary buy the particular radio that was physically in the window or just one of the same type? The interpretation of *it* in (3) offers further possibilities: perhaps Mary’s radio, perhaps the radio in the window, perhaps some other window or collection altogether. The point of the example is that there is no requirement for a language user to resolve these details. Processing might not and need not occur to such a fine grain. Similarly, scope ambiguities need not always be resolved in order for the import of an utterance to be clear. From ‘At least one problem preoccupies every politician’, along with ‘John is a politician’, one can infer that ‘At least one problem preoccupies John’. This inference can be made without committing to one or other scope ordering of the original sentence.

These examples suggest that language users often leave details unresolved when there is no requirement for making a decision. The point can be extended beyond reference and quantifier scope resolution. Frisson (2009) surveys experimental work that appears to show a similar effect at work with polysemy. For example, the word ‘school’ can be used in ways that would invite, if required, subtly different dimensions of implementation. Jocelyn walked to the school (emphasis on the physical place); her mother talked to the school (this is manifested by talking to representative members of staff); the school won the match (an action that is implemented by members of the school); and so on. Frisson speaks of these different implementations as interpretations or senses that the word ‘school’ can take on in different uses.

Experimental evidence shows that these interpretations are not accessed in the same way that the differing meanings of homonymous words are accessed. With homonyms (such as ‘coach’: vehicle, trainer), experimental work shows that all the possible meanings are initially

activated, with access largely determined by the frequency of the individual meaning (a low frequency meaning will be activated but might not reach the threshold required for access). Context can reorder access by making a subordinate meaning more accessible. Resolution of one or other readings typically occurs without delay. For example, on hearing ‘The coach argued with the players’, both meanings of ‘coach’ will initially be activated but the vehicle meaning will quickly be dropped. Sometimes the wrong choice will be made, and this delays interpretation further down the line. We might have: ‘The coach argued with the players; it refused to start when they were jumping around inside’.

By contrast, with the ‘school’ examples (and other similar examples), experimental evidence shows that relative frequency of the various possible senses has no effect and no extra processing effort was found when a less frequent sense was intended. While a possible account is that all possible senses are initially activated (independent of frequency), Frisson and Pickering (1999) argue that there is good reason to rule this out, on grounds of being uneconomical (words often have lots of different senses) and of there being no evidence for competition between the senses (which would be expected if all senses were activated and would lead to longer reading times).¹³ Further, if all senses were activated, there would be a big problem with sense selection (far more so, given the number of senses, than with typical homonymy): if the preceding context does not indicate which the appropriate sense is, we would predict that revisions would frequently be needed if a particular sense was in fact selected.

Frisson and Pickering (1999) put forward, therefore, an alternative hypothesis: what is initially activated is a single, semantically underspecified meaning. This meaning is meant to encompass all the semantically related interpretations of a word.

This abstract meaning is supposed to be the same for all established senses of a word, that is, the same underspecified meaning encompasses all semantically related interpretations of a word that are known to the reader. (Frisson, 2009)

Frisson combines this suggestion with the claim that language users make do with ‘good enough’ interpretations, and suggests that language users will often not bother to resolve the interpretation of ‘school’ beyond this initial underspecified abstract meaning. All options are left open, and the actual processing that goes on may be shallow: ‘a GE [good enough] approach seems to be most compatible with the initial activation of rather shallow lexical-semantic information, which is implied in the idea of semantic underspecification’ (Frisson, 2009).

However, while Frisson’s work helps show how the understanding that we immediately pick up from words need not show depth, there are aspects of his approach to underspecification that can be questioned. The ‘school’ type of example is presented in the following way:

- The different possible readings are called ‘interpretations’ or ‘senses’.
- The underspecified meaning is the ‘same’ for all these senses and ‘encompasses’ them all.

This pushes us in the direction of thinking that ‘school’ has a wide array of interpretations (e.g., covering both physical buildings and people), and that all these interpretations are somehow

¹³ On the contrary, lexical decision tasks for words with multiple senses tend to show faster response times than for words with fewer senses.

present in embryo, or otherwise ‘encompassed’, in a single meaning of ‘school’. This single meaning does not itself fully specify which interpretation will be relevant on an occasion.

Frisson himself is not committed to what form this underspecified meaning will take (see Frisson (2009), 121f.).¹⁴ One option is that the meaning underspecifies in the sense of providing an over-rich semantic representation, from which some sort of selection is required for achieving the interpretation required for a particular occasion of use (Pustejovsky, 1995) presents an ‘over-rich’ approach to lexical semantic representation). An over-rich representation contains more information than is required for a given interpretation, and does not in itself specify which aspect of that over-rich representation will be relevant in a given use. A second option is to go in the opposite direction and to think of a highly schematic or impoverished type of meaning that contains some core commonality that holds across all interpretations. Rather than requiring a process of selection, a schematic meaning needs to be ‘filled out’ if we are to get to a specific interpretation. We might think of a single skeleton that can be fleshed out in different ways on different occasions.

While Frisson is non-committal, I wish to focus on the second (‘impoverished’) option. In part this is because I think that simpler, more skeletal, accounts of lexical knowledge are more plausible than accounts that propose more complex types of representation. But also, it is difficult to understand how a rich semantic specification fits Frisson’s general approach. He writes that what is initially accessed in interpretation is “not a full-fledged, specific interpretation”, but rather something that “will serve as the starting point for a more detailed specific interpretation” (Frisson 2009, p. 122). Further, in line with the ‘good enough’ emphasis, what is initially activated is “rather shallow lexical-semantic information” (Frisson 2009, p. 123). It is not clear how a rich lexical entry can be characterized as requiring ‘more detail’ or as being ‘rather shallow’.

However, it is also problematic to understand how an impoverished meaning can in fact play the kind of role that is required by the example that Frisson provides. First, there is no clear sense as to what this impoverished representation that is shared across all interpretations will look like. With the term ‘school’, we seem to be positing an abstraction that captures a commonality shared by physical buildings, sports teams, administrators, institutions, and other possible nuances of interpretation. It is hard to imagine what this commonality could be.

Carston (2016), criticizing ‘schema’ approaches in general, raises further concerns. One concern is that the proposed abstraction doesn’t explain the various interpretations but is merely a result of those interpretations. If a new interpretation arises, the abstracted meaning will have to change in some way to accommodate it.¹⁵ Explanation starts with the interpretations, and the proposed meaning is nothing but a by-product of this—it is an explanatory idle wheel. Further, the proposed abstraction does not seem to have any relevance to what children learn when they learn words. A child does not seem to learn the abstraction, but would appear to start with one or other of the various interpretations.

The account of word meaning in terms of analogical cognition provides a response to these concerns, as well as challenging aspects of Frisson’s discussion of the ‘school’ example. Rather than saying that the meaning of ‘school’ is underspecified, or that it encompasses all the interpretations, the claim is that the meaning of ‘school’ is completely specific. That is to say, speakers of English share a specific understanding, an understanding that grasps the structural relation that is describable roughly as ‘institution that gives teaching’. Nothing

¹⁴ Thanks to Josie Bowerman and Robyn Carston for pointing this out.

¹⁵ This point is noted by Frisson and Pickering themselves: ‘The underspecified meaning is only applicable for established senses. Because the underspecified meaning is an abstraction over the features of specific senses, a novel interpretation of a word cannot be captured by the underspecified meaning’ (Frisson and Pickering, 2001).

further is specified by the understanding; but in its own terms this understanding is specific, and by hypothesis this specific notion of school occurs as a clearly articulated constituent of thought that is triggered when the word is used.

More detailed consideration of what the various scenarios involve might lead to a focus on specific aspects of the institution; but this does not undermine the claim that all that the word is responsible for is the structural relation that is triggered in thought. More detailed consideration might lead us to conclude that Jocelyn's mother spoke to a member of staff. But this doesn't need to be encompassed by the meaning of 'school', in the sense of requiring a meaning that is somehow general across this and other implementations. Sainsbury (2001) argues that there is a tendency for the claim that there are distinct ways in which one and the same meaning could be true to be confused for the claim that there are distinct meanings. The meaning of the word need not be responsible for any of the (endless) possible details concerning how the scenarios are implemented. In particular, notice that we may sometimes deliberately speak in terms (for example) of having contacted 'the school' when we do not wish to focus on the specific person through whom this contact was made. Our point may simply be one of indicating that the institution was informed. And we may just have used a loudspeaker. On the current hypothesis, the word 'school' is not itself responsible for any particular implementation.

This revised approach indicates that the 'good enough' reading is not to be understood in terms of a mysterious abstraction, but in terms of a quite straightforward understanding that we are talking about a school, viz., an institution that provides education. This will be good enough for many purposes, and if more information is needed it will be elicited (e.g., 'Who did you speak to?').

We can, therefore, make a distinction between plausible and implausible abstractions. The plausible abstraction relates to the kind of structural relation that Gentner has studied. The implausible abstraction is the type of abstraction that is meant to provide a core commonality that embraces the potentially many different accounts of precisely which part of the institution was directly relevant in a given instance. Notice in particular that the range of interpretations that Frisson has in mind relate to particular features of schools, hence *not* to the overall relational perspective. The focus is more on a kind of general knowledge about schools (such as who one speaks to when one contacts a school), but it is just this kind of general knowledge that the account based on analogical cognition indicates is not relevant to word understanding.

The account provides a response to Carston's idle wheel criticism. Rather than being a putative abstraction that results (perhaps implausibly) from a wide and possibly growing array of interpretations, the proposal is for a psychologically realistic level of understanding that plays a central explanatory role in the judgements that are in play when words are used. Carston (2016) supports her idle wheel criticism by arguing that when children first learn a word, what they learn are precisely the senses that are grasped in communication, not the proposed schematic abstractions. But we have seen that there is good evidence, even from a very young age, that word use induces a focus on structural relations—these being plausible abstractions—and that subsequent use of a word triggers the corresponding understanding. Further, the account is not suggesting that word understanding starts at a level that is abstracted away from normal everyday thought. On the contrary, the triggered element of thought can be exactly what is being conveyed. In Carston's terms, the account is fully conceptual.

By appealing to specific, highly accessible, and psychologically realistic understandings of individual words, we have a plausible explanation for the immediacy with which understandings are triggered on hearing language. Details of the features of items or situations are left unspecified, but language users make do with 'good enough' representations. Not all

aspects of interpretation, such as particular implementations for the school scenarios, need to be accounted for in a theory of word understanding.¹⁶

7 A contribution to cognition

The role of language is sometimes envisaged as a medium through which the conceptual realm can be accessed and communicated. A. Clark has described this as the ‘translation view of language’, in which encountered language ‘merely serves to activate the complexes of internal states or representations that are the real cognitive workhorses’ (Clark, 2006). Once the translation has been accomplished, the linguistic form can be thrown away, leaving just the ‘essence’, a meaning that has been ‘fully extracted and rendered in some alternative inner format’.

Gentner’s work supports an alternative view: words are tools that contribute aspects of cognition that are either not found or are under-represented at the more basic biological level of thought: ‘One function of language may be to augment natural modes of cognition with an alternative representational scheme that permits abstract cognition’ (Gentner, 2003). In particular, words may contribute by providing memorizable and context *insensitive* building blocks (as Clark (2006) puts it), and this is possible by virtue of the shift to analogical cognition that the use of linguistic form facilitates. From this perspective, to talk of words as ‘pointers’ to complex arrays of conceptual material is misguided. Rather than pointing somewhere, words trigger specific structural understandings, which help take our thought to a different level.

Much experimental work shows that if people, in a given situation, are asked to think of situations that are similar to the one they are in, their decision ‘is typically driven largely by surface similarities, such as similar objects and contexts, rather than by similarities in relational structure’ (Gentner & Smith, 2013). In other words, the kind of similarity that promotes memory retrieval tends to contrast with the kind of retrieval that is facilitated by language.

The claim that ... language aids analogical retrieval is important, because analogical retrieval is generally quite poor. People routinely fail to be reminded of past experiences that are relationally similar to current experiences.

(Gentner & Smith, 2013)

This suggests that judgements of similarity, when not guided by use of a given word, will not group situations in the way a typical word of language would. The word ‘open’ groups together widely disparate items and situations that do not share surface similarities (open books, bottles, rooms...). By contrast, a judgement of similarity, not guided by the word ‘open’, would, in general, be expected to group an open bottle with a surface similar item such as a vase, an open book with a neat sheaf of papers, and so on. This is a problem for accounts of words, such as Recanati’s sketch of ‘meaning eliminativism’ (2004), that appeal to judgements of similarity—between a situation in which a word has already been used and a novel situation—as the basis

¹⁶ Nathan Klinedinst has pointed out to me that if we take ‘school’ to have a single meaning of the type I suggest, we are implicitly loosening selectional restrictions for other predicates. For example, we say ‘The school is made of wood’, though an institution that provides teaching, being partially abstract, cannot itself be the sort of thing that is wooden. Further, we have the problem of explaining why certain inferences seem to be blocked. From ‘The school is made of wood’, and ‘The school won the match’, we don’t want to infer ‘Something made of wood won the match’. Perhaps a partial explanation is that when further details or inferences are elicited, aspects of the ‘good enough’ interpretation can get challenged and thereby need to be modified. But these modifications are only required when circumstances demand it.

for explaining when a word is re-used. Arguably, such accounts will make the wrong predictions about word usage. Rather than a term being used of situations that we are naturally inclined to regard as similar, it seems as if the term itself plays a central role in facilitating and directing the relational focus that is manifested in the typical array of uses found with words.

8 A way forward?

In summary, human cognition arises in part ‘from learned symbol systems that facilitate the apprehension of relational structure’ (Gentner, 2003). Language provides a structure that is immediately apprehended and that interacts with our cognitive abilities, resulting in the kind of thought that humans are capable of. Each individual word plays a small part, contributing a specific piece of structure.

It remains to be seen whether this approach is consistent with the various phenomena (grouped under headings such as polysemy, metonymy, metaphor) that have been studied by theorists. By looking to a basic underlying understanding of a word, the hypothesis has considered only a limited, albeit important, part of the puzzle. In practice, words appear together in syntactic structures, and this contributes another major component. But in any explanation, it is important to have some idea of the division of labour (what different elements of language can and cannot contribute), and of the kinds of representations that need explaining. An approach to words based on analogical cognition helps to provide some insight into these issues.

References

- Bach, T. (2012). Analogical cognition: applications in epistemology and the philosophy of mind and language. *Philosophy Compass*, 7(5), 348–360.
- Borer, H. (2005). *Structuring sense volume 1: In name only*. Oxford: Oxford University Press.
- Borg, E. (2012). *Pursuing meaning*. Oxford: Oxford University Press.
- Carston, R. (2002). *Thoughts and utterances*. Oxford: Blackwell.
- Carston, R. (2012). Word meaning and concept expressed. *The Linguistic Review*, 29(4), 607–623.
- Carston, R. (2013). Word meaning, what is said and explicature. In C. Penco & F. Domaneschi (Eds.), *What is said and what is not*. CSLI Publications: Stanford, CA, USA.
- Carston, R. (2016). The heterogeneity of procedural meaning. *Lingua*, 175–176.
- Christie, S. & Gentner, D. (2014). Language helps children succeed on a classic analogy task. *Cognitive Science*, 38, 383–397.
- Clark, A. (2006). Material symbols. *Philosophical Psychology*, 19(3), 291–307.
- Ferreira, F., Bailey, K., & Ferraro, V. (2002). Good-enough representations in language comprehension. *Current Directions in Psychological Science*, 11, 11–15.
- Fodor, J. (1998). *Concepts*. Oxford: Clarendon Press.
- Fodor, J. (2008). *LOT 2*. Oxford: Oxford University Press.
- Fricker, E. (2003). Understanding and knowledge of what is said. In A. Barber (Ed.) *Epistemology of language* (pp. 325–366). Oxford: Oxford University Press.
- Frisson, S. (2009). Semantic underspecification in language processing. *Language and Linguistics Compass*, 3(1), 111–127.
- Frisson, S. & Pickering, M. (1999). The processing of metonymy: evidence from eye movements. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 25(6), 1366–1383.
- Gentner, D. (2003). Why we’re so smart. In D. Gentner & S. Goldin-Meadow (Eds.), *Language in mind: advances in the study of language and thought* (pp. 195–235). Cambridge, MA: MIT Press.
- Gentner, D. (2005). The development of relational category knowledge. In L. Gershkoff-Stowe & D. H. Rakison (Eds.), *Building object categories in developmental time* (pp. 245–275). Hillsdale, NJ: Erlbaum.
- Gentner, D. (2010). Bootstrapping the mind: analogical processes and symbol systems. *Cognitive Science*, 34, 752–775.

- Gentner, D. & Christie, S. (2010). Mutual bootstrapping between language and analogical processing. *Language and Cognition*, 2(2), 261–283.
- Gentner, D. & Kurtz, K. (2005). Relational categories. In W. Ahn, R. Goldstone, B. Love, A. Markman, & P. Wolff (Eds.), *Categorization inside and outside the laboratory* (pp. 151–175). Washington, DC: APA.
- Gentner, D. & Markman, A. (1997). Structure mapping in analogy and similarity. *American Psychologist*, 52(1), 45–56.
- Gentner, D. & Smith, L. (2013). Analogical learning and reasoning. In D. Reisberg (Ed.), *The Oxford Handbook of Cognitive Psychology* (pp. 668–681). New York: Oxford University Press.
- Hoyningen-Huene, P. (2015). A note on the concept of game (or rather *Spiel*). In G. Betz, D. Koppelberg, D. Löwenstein, & A. Wehofsits (Eds.), *Weiter denken – über Philosophie, Wissenschaft und Religion* (pp. 205–210). Berlin: de Gruyter.
- Keil, F. & Kominsky, J. (2015). Grounding concepts. In E. Margolis & S. Laurence (Eds.), *The conceptual mind* (pp. 677–692). MIT Press.
- Loewenstein, J. & Gentner, D. (2005). Relational language and the development of relational mapping. *Cognitive Psychology*, 50, 315–353.
- Penn, D., Holyoak, K., & Povinelli, D. (2008). Darwin’s mistake: explaining the discontinuity between human and nonhuman minds. *Brain and Behavioral Sciences*, 31, 109–178.
- Pettit, D. (2010). On the epistemology and psychology of speech comprehension. *The Baltic International Yearbook of Cognition, Logic and Communication* (Vol. 5, *Meaning, Understanding and Knowledge*, pp. 1–43).
- Pietroski, P. (2005). Meaning before truth. In G. Preyer & G. Peter (Eds.), *Contextualism in philosophy* (pp. 255–302). Oxford: Clarendon Press.
- Prior, W. (1983). The concept of paradigm in Plato’s theory of Forms. *Apeiron*, 17(1), 33–42.
- Pritchard, T. (forthcoming). Is there something in common? Forms and the theory of word meaning. *European Journal of Philosophy*.
- Pustejovsky, J. (1995). *The generative lexicon*. Cambridge, MA: MIT Press.
- Recanati, F. (2004). *Literal meaning*. Cambridge University Press.
- Rayo, A. (2014). A plea for semantic localism. *Nous*, 47(4), 647–679.
- Sainsbury, M. (2001). Two ways to smoke a cigarette. *Ratio*, 14(4), 386–406.
- Sanford, A. & Sturt, P. (2002). Depth of processing in language comprehension: not noticing the evidence. *Trends in Cognitive Sciences*, 6(9), 382–386.
- Sperber, D. & Wilson, D. (1998). The mapping between the mental and the public lexicon. In P. Carruthers & J. Boucher (Eds.), *Language and thought: interdisciplinary themes* (pp. 184–200). Cambridge: CUP.
- Travis, C. (2000). *Unshadowed thought*. Cambridge, MA: Harvard University Press.
- Wilson, D. & Carston, R. (2007). A unitary approach to lexical pragmatics: relevance, inference, and ad hoc concepts. In N. Burton-Roberts (Ed.), *Pragmatics* (pp. 230–259). Palgrave: London.
- Wittgenstein, L. (1953). *Philosophical investigations*. Oxford: Blackwell.