

# **Literal and metaphorical meaning: In search of a lost distinction**

*Nicholas Allott, University of Oslo, and Mark Textor, KCL*

Draft: Feb 2022. Please send any comments to [nicholas.allott@gmail.com](mailto:nicholas.allott@gmail.com)

## **Abstract**

In recent years the distinction between literal/non-literal truth and meaning has come under attack. On a traditional conception, figurative uses are ones in which the meaning expressed deviates from the literal meaning. But recent work on lexical modulation and polysemy shows that meaning deviation is ubiquitous, even in cases of literal use. In this paper we argue that the literal-figurative distinction should be preserved and not eliminated. Even if utterance literalness does not pull its weight in a theory of language processing, the distinction between literal and figurative is practically and theoretically important. We show that one can give up the idea that there is a standing meaning that is expressed in utterances that are literally true but keep the core of the folk-theoretic notion of the literal/non-literal distinction. Our claim is that literal uses of a word are made with the intention to conform to an established practice of use, while metaphorical uses do not so conform, but depend on this pre-existing practice. We show that our account can deal with data that are problematic for other theories.

## **1. Introduction: The Literal/Non-Literal Distinction in the Firing Line**

There are distinctions that are anchored in common sense and clarified by philosophy and science that seem to be undermined by progress. A good example is the distinction between different senses: sight, touch etc. Common sense distinguishes between five senses. But psychological and neurophysiological theorizing seems not to vindicate this distinction.<sup>1</sup> Hence, we need either to find a new way to vindicate the common sense distinction or give it up.

In recent years a similar issue has arisen for the distinction between literal and non-literal use. We are generally good at distinguishing between literal and non-literal meaning and truth: ‘The brightest object visible from earth is the sun’ is literally true,

---

<sup>1</sup> For an overview and discussion see Nudds 2004.

‘Juliet is the sun’ is not. Recanati (2004: 68) submits that we have a folk-theoretic notion of the literal/non-literal distinction:

In the ordinary sense of the term, non-literal meaning contrasts with normal meaning. Non-literal meaning is special, it involves a form of deviance or departure from the norm; a form of deviance or departure which must be transparent to the language users. (Recanati 2004: 81)

This folk-theoretic notion sees the literal as basic or standard and the non-literal as a deviation from the literal. To anticipate: in the positive part of this paper we will support and defend the *core* of the folk-theoretic notion of the literal/non-literal distinction.

The way the folk mark the distinction between the literal and the non-literal needs support because it seems to rely on a conception of meaning that has come under pressure in the last twenty years. This conception assumes that literal meaning pertains to the expression-*type* and is determined prior to and independently of contexts of use<sup>2</sup>. The non-literal meaning of a word-type *w* deviates from the literal meaning and is not fixed in the same way.<sup>3</sup> If we already have in place a notion of literal meaning for a sentence type, it is tempting to say that an utterance of a sentence *s* is literally true if the utterance expresses the literal meaning of the sentence and things are as the literal meaning of the sentence represents them. When does an assertoric utterance of a declarative sentence express its literal meaning? If a speaker utters the sentence intending to express its literal meaning.<sup>4</sup> In this case the question of literal truth/falsity arises for the utterance. Otherwise the utterance is not literally true. It may be metaphorically or approximately true etc.

Now, many philosophers of language and linguists have given good arguments for the conclusion that deviation from the meaning of an expression-type, if there is

---

<sup>2</sup> In our view, the linguistic semantics of a word is determined by a combination of conventions and cognitive abilities. For details see (omitted), and below.

<sup>3</sup> For the deviation view see White 2001: 24, 32-3; Recanati 2004: 81; Stern 2006: 250.

<sup>4</sup> On a Gricean view of speaker meaning, this amounts to intending that its meaning is entertained or accepted by their audience. We doubt that claim, but do not discuss alternatives here.

such a thing, happens more often than not. To put it somewhat paradoxically: deviation from the norm is the normal case, not the exception. Work on lexical modulation in linguistic pragmatics and on polysemy in lexical semantics converges on this conclusion. Consider a simple example. When you say ‘John cut the cake’, the meaning of ‘cut’ in your utterance is not the meaning, if any, of ‘cut’ independently of a particular utterance. The meaning of ‘cut’, if there is one, is so to say fine-tuned to mean cutting appropriate to cake and the particular circumstances (Searle 1980). However, while the meaning ‘cut’ has in this particular utterance is distinct from the meaning of the word type ‘cut’ it is certainly not non-literal.

Some may argue that there is no occasion-specific meaning here, that ‘cut’ simply has its literal standing meaning, even though a hearer of the utterance typically forms a more specific view of the manner of cutting<sup>5</sup>. Such scepticism is harder to maintain in the face of examples like ‘John’s new book is heavy’ and ‘John’s new book is available on Kindle’, where ‘book’ has a concrete sense (book object, as it were) in the former, but in the latter means something like book content.<sup>6</sup> These different senses of polysemous words are components of the content of the speaker’s assertions (for evidence see Collins 2017: 679).

There is evidence from psycholinguistic experiments on polysemy (e.g. Frisson 2009; 2015) that the stable, linguistically encoded meanings of polysemous words are neutral between word token meanings and not equal to any of them. We cannot explain the experiments in detail here<sup>7</sup>. Briefly, they show that the difference between homonyms (e.g. *bat* [rodent] and *bat* [sports]) and polysemes (e.g. different senses of the noun *book*) is cognitively real: they are processed differently. In particular, in processing polysemes, there’s no bias towards the most frequent or otherwise dominant sense (Frisson 2009), while it has long been known that there is such a bias for homonyms. Also, polysemous senses prime each other, while homonyms inhibit each other. The conclusion Frisson draws is that the standing linguistic meaning of a polysemous word is something that is neutral between its senses and not equal to any of them. So the meaning of ‘cut’ in a competent speaker’s mental lexicon is neutral between the senses that can be expressed when a speaker says ‘cut grass’, ‘cut the

---

<sup>5</sup> Fodor (1998: 53–4) pursues this strategy vs. the apparent polysemy of ‘keep’.

<sup>6</sup> See also Searle 1983: 146; Vicente 2019: 921–2.

<sup>7</sup> For a good summary see Quilty-Dunn 2020: 8–12.

cake', cut hair' etc., and similarly for the abstract and concrete senses of 'book'. On this view, all uses of polysemous words (i.e. of most content words<sup>8</sup>) are 'deviant' in the sense discussed above: the word token meaning is *never* identical to the word type meaning.

If meaning deviation is ubiquitous, it seems that one can either (a) hold on to the distinction between literal/non-literal meaning, but disentangle it from the deviation explanation of it or (b) accept that the deviation model is the only adequate account of the distinction and therefore abandon the literal/non-literal distinction.

White (1996; 2001) pursues (a). He writes:

Once we abandon the idea that the literal meanings of words are regulated by convention, and attend to the ways in which words used literally can continually change their meanings from context to context, then talk of a "deviant" meaning loses all purchase. (White 2001: 58)

But since he believes that there is a distinction between literal and metaphorical truth, he seeks to explain it independently of the deviation model of non-literalness.

In contrast, Wilson and Sperber lean towards (b):

The interpretation of every utterance involves a process of meaning construction, which is the same whether the result is an enriched, loosened, enriched-and-loosened, or literal interpretation. Yes, literalness can be defined, or at least characterized, in terms of a prototype, but no, verbal understanding does not involve paying any attention to literal meaning, let alone to literal use. There is no theoretical basis for sharpening our characterization of literalness. (Wilson & Sperber 2002: 624)

On this view, the meaning of (almost) every utterance of a content word is constructed on the fly and there is no presumption that the standing (or 'linguistically encoded') meaning of the word-type will be expressed. We may have a folk-theoretic distinction between literal and figurative use, but in the development of theoretical views of meaning and understanding it can and should be abandoned.

---

<sup>8</sup> 80% on one estimate: Rodd 2020: 411.

Which way should one go? In this paper we will argue that the literal-figurative distinction should be preserved and not eliminated. Even if utterance literalness does not pull its weight in a theory of language processing, the distinction between literal and figurative is practically and theoretically important (section 3). Unlike White, we will not give up the deviation model: our goal here is to save it. But we will give up the idea that there is a standing meaning that is expressed in utterances that are literally true.

## **2. Lexical Pragmatics and the Literal/Non-Literal Distinction**

Let us start by looking in more detail at an influential view of lexical pragmatics that is in tension with the literal/non-literal distinction. Sperber and Wilson assume that monomorphemic words encode atomic concepts that in turn provide access to encyclopaedic information about the objects in their extension (1986: 87; 91–2). According to the account of lexical pragmatics they and Robyn Carston have developed, in interpretation hearers construct senses by selecting encyclopaedic properties (Carston 1997; Sperber & Wilson 1998; Wilson & Carston 2007).<sup>9</sup> These are selected in order of accessibility (which is affected by context) and the candidate sense (*‘ad hoc concept’*) depends on which are selected. The account claims that *ad hoc* concepts have extensions that differ from the lexicalised concept: they are broader or narrower. (We return to this point below.) This process occurs in mutual adjustment with other pragmatic processes including disambiguation, assignment of referents to indexicals and generation of implicatures. The result is a candidate interpretation of the utterance, which is accepted by the hearer’s pragmatic system as the speaker-intended one if it satisfies expectations of relevance in the context. (For helpful examples, see Sperber & Wilson 2002: 19–20; Carston & Powell 2006: 345–6)

According to this account, the generation of *ad hoc* concepts is operative in almost all utterances (Carston & Powell 2006: 345; Wilson & Carston 2007: 231). Consider typical utterances of the following three sentences:

---

<sup>9</sup> For Wilson and Carston ‘encyclopaedic property’ is a shorthand for proposition that has an encyclopaedic property as a constituent: ‘each property should be seen as a constituent of a complete proposition’ (2007: 256, n. 22).

- (1) The surgeon opened John's mouth. (His jaw was broken and had been wired together) (Carston & Powell 2006: 344)
- (2) Paul Pogba was born in Paris. (He was in fact born in a town well outside of the municipal boundary of Paris which is nowadays regarded as a suburb of Paris)
- (3) Paul Pogba is no saint.

If there is a standard context-independent meaning of the underlined words, it is not the meaning put forth in the utterances we are concerned with. But, according to common sense, it is only (3) that is a clear case of figurative speech. Similarly, as we have seen 'cut' seems to have different – but equally literal – senses in typical utterances of 'cut hair', 'cut cake', 'cut grass' and so on (Searle, 1980: pp. 222–3; cf. Searle 1983: 145–6 on 'open').

One conclusion one might draw from these observations is that, given the omnipresence of deviant occasion-sensitive meaning, distinguishing between literal and figurative meaning is pointless. If everything is a deviation, nothing is: the distinction between deviation and norm loses its point and purpose and can no longer ground the literal/non-literal distinction. Hence, we should give up the later distinction. Or so the revisionists argue.

But can't the distinction between literal and figurative consist in a difference in the ways in which one 'constructs' the deviant meaning? Not on this view, for the same method of arriving at the occasion sensitive meaning is operative in arriving at the occasion sensitive meaning of utterances of (1) to (3):

[F]ully unified accounts reject the traditional distinction between literal and figurative meaning and claim that approximation, hyperbole and metaphor are not distinct natural kinds, requiring different interpretive mechanisms, but involve exactly the same interpretive processes as are used for ordinary, literal utterances. (Wilson & Carston 2007: 231)

Similar conclusions have been widely reached in experimental pragmatics. The key finding here is that, "with enough background information, a novel metaphoric phrase can be read as routinely as a non-metaphoric one." (Noveck 2018: 159). This is evidence against the literal-first hypothesis: that in processing a metaphorical

utterance the addressee must first entertain a literal interpretation, then reject it and construct the correct metaphorical interpretation. It also suggests a more general conclusion: metaphorical utterances are processed the same way as literal ones. As Noveck says, reviewing the experimental literature:

It would soon become commonplace to find articles whose introductions resembled the following: “Revered scholars have thought that there is something special, unique, or exceptional about metaphor and now we know that this is not the case.” (Noveck 2018: 160)<sup>10</sup>

Likewise in lexical semantics, as witness the scare quotes in the title of a recent paper, *‘Figurative’ uses of verbs and grammar* (McNally and Spalek 2019). This paper (which we return to below) treats the literal/figurative distinction as intuitively obvious in some cases, contrasting e.g. ‘cut the cake’ and ‘cut the deficit’, but does not attempt to provide any theoretical basis for this distinction. Rather, it argues that ‘literal’ and ‘figurative’ uses are similar, as we shall see.

### **3. Why the Literal-Figurative Distinction is Needed**

So there are reasons that undermine a view of the literal/non-literal distinction. Should we, then, give up this distinction? No, the distinction between literal and figurative uses is one which we make and want to continue making when we progress to scientific thought. Consider an example. You are interested in finding out what kind of thing the mind is. A cognitive scientist who has gathered evidence and has formed a theory tells you:

(4) The mind is a computer.

You are surprised. Minds are not artefacts and don’t seem to be programmable, have a central processing unit etc. It’s reasonable to ask: Is the mind literally a computer? If the answer is ‘No’, the cognitive scientist has no theory of the mind yet. Why not? If it is only metaphorically true that the mind is a computer, we have only (at best)

---

<sup>10</sup> A prominent example is Giora 2008.

started saying what the mind is.<sup>11</sup> Roughly speaking, the cognitive scientist draws our attention to something that is in certain respects analogous to the mind. But drawing our attention to the right kind of model for the development of a theory is not having developed and proposed a theory. It is a starting point, not the end point. Whether an utterance is put forth as literally or metaphorically (more broadly: figuratively) true matters for assessing the status of a proposal: is it a heuristic for finding an explanation or is it the explanation we seek? Here we find ourselves in agreement with Colyvan (2010: 299): a metaphor is only a proxy for an explanation that can be given without it. This thought about metaphor nicely explains some debates in science and philosophy. For example, Chalmers and Clark argue that “My notebook is an extension/part of the machinery of my mind” is ‘quite literally’ true (Chalmers 2008, xxvi). Fodor (2007) denies that it is. If my notebook is only figuratively part of the machinery of my mind, the statement has less value for cognitive science:

If minds don’t literally have parts, how can cognitive science literally endorse the claim that they do? That Juliet is the sun is, perhaps, figuratively true; but since it is only figuratively true, it’s of no astronomical interest.

Fodor’s example is unfair: that Juliet is metaphorically the sun is indeed of little astronomical interest. But that minds metaphorically speaking have parts, among them, notebooks, could be of interest to cognitive science. Still, the interest is only an interest in a ‘line of thought’ or heuristic, not an explanation that can be confirmed or refuted.

Does this imply that only literally true statements can constitute a satisfactory theory? For a while philosophers believed so. Quine (1981: 189) suggests that at least “in the inner stretches of science” tropes are cleared away. But Yablo (1998: 250-3) and others have argued that some metaphors are representationally essential. Camp (2006), for example, claims that the study of the mind at present requires the metaphor of memory retrieval as ‘access’ to a ‘computer file’. Yablo’s view is that in

---

<sup>11</sup> The claim that the mind is a computer has been criticised as a bad metaphor (e.g. Epstein 2016). Two very different responses are possible. Camp (2006) argues that it’s a good metaphor, essential in our current state of ignorance, while Lande (2019) argues that it’s literally true. There’s a *prima facie* serious debate here, and that clashes with the claim that the literal/figurative distinction is ill-founded.



some cases a metaphor cannot be eliminated from the theory without weakening it; Camp's that as progress is made, the metaphor may be dispensable, replaced by a literal concept. Whatever the correct resolution of these debates, it's clear that in science the distinction between a literally true/false *statement* – an assertoric utterance for which the question of truth/falsity arises – and a figurative claim plays an important role.

There are many other areas of life where the literal/figurative distinction matters. The law is one. Drafters of statutes and contracts avoid non-literal expressions, except highly conventionalised ones, with the aim of avoiding unclarity (Tiersma 2001: 454; Lewison 2015: para 15). As with science, there are nuances. Metaphors like 'bite at the apple' are used as memorable shorthands for established legal principles in some other types of legal text, including judgments and written arguments (Richard 2014). But "[e]xperience shows that rules formulated in terms of metaphors always cause trouble when it comes to their interpretation" (Lord Hoffmann in *Serco Ltd v Lawson*, quoted in Lewison 2015: para 10).

The literal/figurative distinction is also important in semantics and philosophy of language, where theoretical claims are often explicitly limited to literal uses of expressions. For example it's not uncommon for non-literal uses to be excluded in work on the meaning of slurring terms. Jeshion writes, "as in other semantic theorizing, we must carefully separate literal uses of slurring terms from non-literal uses" (2013: 234) and she treats what she calls 'non-weapon' uses of slurs as non-literal (p. 247), allowing her to claim that "slurs' capacity to derogate is conventional; it is a feature of, a power of, the words themselves" (p. 232). Similarly, Belleri (2020: 9) argues that all 'genuine uses' of slurs are derogatory, where those are restricted to literal uses.

In recent semantic theories that claim that meanings of sentences are not truth-conditional but are constraints on what they can be used to express, non-literal uses are set aside. For example, Harris writes:

Roughly, the semantic value of an expression  $\phi$  is just what a competent speaker can know about what someone would be saying in uttering  $\phi$ ,

*assuming they were speaking literally*, but without any knowledge about the context or the speaker's intentions. (2020: 4, our italics)<sup>12</sup>

Clearly, this approach presupposes the literal/figurative distinction. (It also fits naturally with our project here: setting out a deviation account of the distinction that does not rely on a view of word meanings as context-independent concepts.)

Given the importance of the literal/figurative distinction in science, law and semantics it should be preserved and clarified if possible.

#### **4. Unsuccessful attempts to save the literal-figurative distinction**

In this section we review three proposed criteria for a literal-figurative distinction compatible with ubiquitous polysemy and lexical modulation. We argue that all three fail.

##### **4.1 Recanati's Rescue Manoeuvre**

In his *Literal Meaning* Recanati sets out to show that “the ordinary, folk-theoretic notion of non-literal use can be rescued” (2004: 68) from the claws of Sperber and Wilson's eliminativism. Recanati (2004: 74) agrees with them that lexical modulation is universal and essential for communication. So how do we rescue the ordinary distinction between the literal and figurative?

Recanati claims that non-literalness is ‘transparent’ to language users (2004: 75). In a non-literal use of an expression there is a felt discrepancy between conventional meaning and the utterance meaning. Consider someone saying:

(5) The ATM swallowed my credit card, chewed it up and spat it out.  
(Recanati 2004: 77 fn 11)

Here someone who understands the utterance fully is supposed to feel a ‘tension’ between the standard meaning of ‘ATM’ and the predicates applied to its referent: “In the scene the ATM is both a machine and a human being – hence the feeling of discrepancy. The discrepancy [...] is a tension between two aspects or components of the primary interpretation.” (2004: 79–80)

---

<sup>12</sup> The restriction of constraint-based semantics to literal meaning is also made by Collins (2017: 691 n9).

In contrast, when I say “Paul Pogba is from Paris” there is supposed to be no felt tension. Hence, felt tension is supposed to make for figurative use.

Take a conceptual falsehood – an utterance of (6):

(6) The square Meinong thought of is round.

Squares are not round and if you possess the concepts of square and roundness you cannot fail to know this. Hence, you will feel that just as ATMs cannot swallow credit cards, squares cannot be round. There is a felt discrepancy. But the utterance is not metaphorically or non-literally true. It is literally (and conceptually) false. Similarly, an utterance of ‘Prime numbers are green’ is literally false, but comes with the feeling of discrepancy. A feeling of discrepancy is not sufficient for figurative use, then.

It is also not necessary. There are metaphorical utterances where there is – or need be – no felt discrepancy. Consider again an utterance of (3). It is a clear case of non-literal truth. But there is no felt discrepancy because it is denied that Pogba, the French footballer, is a saint. There are also positive metaphorical utterances (i.e. not denials) with no feeling of discrepancy. Recall the generalisation from experimental work on metaphor processing: where there is a good fit with context, metaphors can be understood as rapidly and easily as literal utterances. Carston (2010) notes that metaphors can differ in a number of respects, including how creative they are, whether they evoke imagery, whether they are spontaneous or highly wrought and literary, and whether they are ‘extended’. She suggests that there are two different modes of processing: “a quick, local, on-line meaning-adjustment process and a slower, more global appraisal of the literal meaning of the whole.” (p. 297). Whether or not that proposal is right, it is highly plausible that in typical processing of a metaphor that is easy in context, not image-evoking, and not part of an sequence of related metaphors, there is no special phenomenology. At the least, the burden of proof falls on the claim that there always is.

We can see two responses for Recanati, one of which he actually makes in introducing his criterion. For Recanati not all metaphors are figurative: only conspicuous ones are. This view has the strange consequence that “The ATM swallowed my credit card” is not figurative, in contrast to (5), which is figurative because adding “chewed it up and spat it out” “belabour[s] the metaphor so as to make it conspicuous” (Recanati 2004: 77 fn 11). So Recanati can deal with the

objection that there are metaphors with no feeling of discrepancy, but only by classifying them as non-figurative! His criterion puts the literal-figurative distinction in the wrong place.

Here's a better response that Recanati could make: felt discrepancy need not always be present, but can be brought to mind whenever there is a metaphor, just as grammaticality intuitions often require some introspection<sup>13</sup>. However, this claim lacks any empirical support, as far as we know, and of course would not help with the sufficiency problem, namely that there are literal utterances that have felt discrepancy.

In sum: We think the prospects are dim for either version of the felt discrepancy view, given there are many ways in which literal utterances can strike language users as deviant either immediately or on reflection and the experimentally supported fact that for normal language users many metaphors do not stand out from other use. If we want to retain the distinction between figurative and nonfigurative meaning and truth, felt discrepancy won't do the job. Still, there is a grain of truth here. Figurative uses are special. We offer an explanation in section 5.

## 4.2 Metaphor as 'radical' broadening

What Wilson and Carston call a 'fully unified' relevance-theoretic account of lexical pragmatics (see section 2 above) would see no qualitative difference in processing between literal and figurative utterances. However, there have been two suggestions by relevance theorists about how one could accommodate the common-sense distinction between one important kind of figurative utterance, namely metaphorical utterances, and literal talk.

Wilson & Carston suggest that metaphor is *radical* broadening:

Within the fully unified account we are proposing, metaphor is seen as a still *more radical variety of broadening* than hyperbole, involving a *greater departure* from the encoded meaning.' (Wilson & Carston 2007: 235, our

---

<sup>13</sup> Cf. the notions of occurrence and availability in work on consciousness. Recanati glosses 'available' as "consciously accessible" (2004: 42) suggesting that he has this distinction in mind. On introspection and linguistic intuitions see Gross 2021: 546: "it can take time and effort to get a reading of a sentence, and reflection and perhaps conscious inference to arrive at a truth-judgment".

emphasis)

What is the measure of ‘greater departure from’ encoded meaning? The examples discussed are all cases in which the extension of the word under consideration becomes more inclusive. We would need a metric to assess *distance from* the encoded meaning or to say in some other way what is a radical broadening. But we don’t get either of these.

Our view is that we can’t have one. The claim here is that it’s in some quantitative sense more of a stretch to extend the meaning of ‘flat’ (for example) so that it applies to a feeling than to extend it so that it applies to Holland. But it seems to be a different *kind* of deviation. We return to this point below.

### 4.3 Metaphor as broadening plus narrowing

Carston and Wearing (2011) suggest that metaphor is broadening *plus* narrowing. They claim that this differentiates it from loose use and hyperbole, which involve only broadening, and strictly literal use (if it exists) where there is neither. Consider (3) again. The *ad hoc* concept expressed by the speaker’s use of the word ‘saint’ here (roughly: exceptionally virtuous) is broader than the literal meaning: there are people who have not been canonised by a church who are this virtuous. It is also narrower: not every literal saint is so well-behaved. This looks promising, but we think it can’t be the right criterion, as we now explain.

A serious problem for views that make broadening and/or narrowing criterial is that encoded and metaphorical meaning, *if any*, can be co-extensional (omitted). Consider:

(7) Caissa is the God of chess.

(8) For many Indian chess fans, Viswanathan Anand is the God of chess.  
(Chessbase India)

Intuitively, the phrase ‘God of chess’ is used literally in (7) and metaphorically in (8). However, they may be co-extensional, since they may both be empty. We take it that it’s not particularly controversial to suggest that in fact there is no deity specifically responsible for chess i.e. ‘the God of chess’ is empty in (7). As for the expression when used in (8), suppose that we find out that all chess grandmasters

have been secretly relying on computer analysis during game play. Then we might agree that no one is good enough at chess to be the God of chess in the (figurative) sense intended.

#### 4.4 Can the last two criteria be salvaged?

It might be argued that we are taking too seriously relevance theory's talk about broadening and narrowing. Perhaps the claim is fundamentally about senses that are modulated, not their extensions: after all, Wilson and Carston describe lexical modulation as "a single [...] process which creates an ad hoc *concept*, or occasion-specific *sense*" (Wilson & Carston 2007: our italics). If so, one should try to re-interpret Wilson and Carston's and Carston and Wearing's claims about broadening and narrowing in terms of word senses<sup>14</sup> or concepts.

Wilson and Carston's claim that metaphor involves a greater departure from the encoded meaning than other kinds of modulation can be understood this way: in metaphorical use a property is expressed which is more different from the encoded sense than in loose use and hyperbole. However, the claim still requires – and lacks – a metric or an explanation of what this greater difference amounts to.<sup>15</sup>

Carston and Wearing's broadening-plus-narrowing view is harder to recast in non-extensional terms. One possibility is to distinguish literal and figurative uses in terms of which 'meaning features' belonging to the encoded sense are deployed in them: broadening of extension can then be understood as a result of 'dropping' a defining feature of the sense, and narrowing could be due to adopting an additional semantic feature. Views of word meanings as (or as including) collections of features are widespread in lexical semantics, and we turn to them next.

#### 4.5 Non-literal use as feature 'dropping'

The evidence about polysemy from Frisson and others (see section 1, above) is compatible with two kinds of theory of word type meaning (Lossius Falkum &

---

<sup>14</sup> We use 'sense' here as in linguistics to mean (roughly) a meaning that a word can be used to express: e.g. as noted above, 'book' has at least two stable senses.

<sup>15</sup> Semantic-space accounts of word meaning (e.g. Rodd 2020) could help provide a metric – distance in high-dimensional semantic space – but would need to explain how a quantitative measure gives rise to a categorical distinction. Such accounts rely on lexical meaning features, for which see the text following this note.

Vicente 2015: 5). These two views share the assumption that word type meanings are (or include) a list or ‘bundle’ of features. (We explain the motivation for this assumption below.)

The first view is that word meanings are overspecified or ‘rich’. Stored word meanings include more features than are typically deployed in any particular use of the word: any word token meaning uses only a subset of the stored features. Some of the stored features may even be incompatible with each other. For example the meaning of ‘book’ may include both the feature <material object> and the feature <intellectual content>.<sup>16</sup> The other view is that word type meanings are underspecified or ‘thin’. On this view, the linguistically encoded features always need to be contextually augmented to arrive at word token meaning.

It’s currently unclear which of these models is correct, and in fact they may both be, but for different classes of words. Vicente (2019) suggests that verbs have ‘thin’ meanings while some nouns have ‘thick’ meanings. Regardless of how current debates play out, if words do indeed have meaning features and not all such features are deployed in every use, then that allows for a criterion for non-literal uses which has been suggested by Carston and Wearing (2011). The idea is that some linguistically encoded word-meaning features are special. If one of these is *not* deployed in the word token meaning, the result is non-literal. For example, ‘cut’ in ‘cut taxes’ is a non-literal use because the feature <affects material object> is not deployed. Similarly, a feature such as <building> or <material object> is not deployed in constructing the word token meaning in ‘John’s job is a prison’, thus this is a metaphorical use.<sup>17</sup>

This criterion looks promising. However, it faces a serious problem, which we first outline, then explain in more detail. It remains controversial that words have

---

<sup>16</sup> Under certain circumstances, apparently incompatible features can be deployed in a single use of a word: e.g. ‘John read and then burned the book’. (See Collins 2017 and Vicente 2019 on these co-predication cases.)

<sup>17</sup> These proposed features of word meaning which are claimed to be necessary for literal use are similar to meaning postulates (Carnap 1947; Fodor, Garrett, Walker and Parkes 1980; Fodor 1981). They may be vulnerable to familiar arguments against the analytic-synthetic distinction and lexical decomposition – but see Rey (2017: §4.4) on I-analyticity.

meaning features.<sup>18</sup> The evidence is stronger for some types of words: essentially, it is much stronger for verbs than other content words where it is very thin (Chomsky 1992: 217). The best-evidenced meaning features of verbs appear to be carried over to metaphorical use (Glanzberg 2008). So there's not much evidence for meaning features that are 'dropped' in figurative use.

A great deal of work in lexical semantics indicates that lexical meaning of verbs divides into two: 'syntacto-semantic' features that are visible to linguistic syntax, plus 'idiosyncratic' content. The features capture thematic role types (whether a verb requires an agent argument, for example) and event structures (e.g. whether a verb is telic – see below for explanation.) Several authors have proposed that this syntacto-semantic skeleton of a word persists in metaphorical uses (Glanzberg 2008; Ramchand 2014; Spalek 2015; McNally & Spalek 2019). We set out the argument using data from Glanzberg. The Italian word normally used to translate English *blush* is *arrossire*. These verbs have very similar meanings, but differ in their lexical aspect. *Blush* is an atelic verb: it means something like "have red cheeks", while *arrossire* is a telic verb: its meaning is closer to "come to have red cheeks" (cf. English *redde*n).

That difference accounts for the following contrast in acceptability:

- (9a) John blushed all day long.  
(9b) \*John arrossí tutto il giorno.

Now metaphorical uses of both verbs are possible, as in the following examples:

- (10a) The sky blushed.  
(10b) Il cielo arrossí.

Crucially, they obey the same constraint as literal uses, as the following examples show:

- (11a) The sky blushed all day long.  
(11b) \*Il cielo arrossí tutto il giorno.

---

<sup>18</sup> E.g. Fodor (1998) argues against lexical decomposition and semantic features, Pietroski (2003) in favour of some analytic truths.



This is not an isolated case. McNally and Spalek show a similar continuity of differences between a number of English and Spanish verbs including English *cut* and its Spanish counterpart *cortar* across literal and figurative uses. In general, they say:

[T]he respective conventional event structures and thematic role types that a given set of counterpart verbs impose on the events they (literally or figuratively) describe will remain consistent within each language for both literal and syntactically productive figurative uses, and potentially vary across languages. (McNally & Spalek 2019)

These findings do not disprove the claim that there are word meaning features whose ‘dropping’ (non-deployment) is necessary for non-literal uses, but they put it under considerable strain. We therefore suggest a different direction.

## 5. Resurrecting the Literal/Figurative Distinction

Our claim is that literal uses of a word<sup>19</sup> are made with the intention to continue a ‘tradition’ of using a word. The central case is that the speaker intends to conform and use the word as other people, in particular its originators, use or used it.<sup>20</sup> A metaphorical use is one which does not so conform, but depends on the pre-existing practice which runs through the literal uses. So there is a distinction between a basis on the one hand and a use that depends on the basis on the other hand and deviates from it. In our view, this kind of distinction between basis and deviation is crucial to any understanding of the distinction between literal and non-literal meaning and truth. The non-literal is derived from and therefore dependent on the literal. Our view is thus a kind of ‘asymmetric dependency’ view. Both literal and metaphorical uses depend on the conventional meaning but the convention does not depend on metaphorical uses.<sup>21</sup>

---

<sup>19</sup> We regard homonyms as different words. We are *not* claiming that all literal uses of the string ‘bank’ contribute to one tradition.

<sup>20</sup> We are influenced here by the view of concepts developed by Schroeter and Schroeter (2016), and the view of words set out in (omitted).

<sup>21</sup> On asymmetric dependency theories of meaning, see Fodor 1987; Rey 2017. Rey also suggests that figurative uses asymmetrically depend on literal uses.

To see that this is plausible, consider this challenge: *Coin a new monomorphemic word and simultaneously use it metaphorically!*<sup>22</sup> Can you do it? No, you can't! Why? Because the literal meaning and use has to be established before you can deviate from it. (One cannot start and subvert a new fashion at the same time).

According to our proposal, there are two ways in which the use of a word can catch on:

Conservative Chain: the word is originated and it is passed on from speaker to speaker where each of the speakers has the intention to conform to the people from which he acquired the word.

In a conservative chain the semantic value of a word may change because of mistakes speakers make. Gareth Evans made this point by means of his well-known 'Madagascar' example.

We learn from Isaac Taylor's book: *Names and their History*, 1898:

"In the case of Madagascar a hearsay report of Malay or Arab sailors misunderstood by Marco Polo ... has had the effect of transferring a corrupt form of the name of a portion of the African mainland to the great African Island." (Evans 1973: 196)

According to Evans' description, the same name 'Madagascar' changed its referent because of Marco Polo's mistake. The important observation for our purposes is that the uses of the name before and after the change of reference are all made with the intention to conform. This thought motivates the idea of a *conservative chain* of uses. An originating use introduces a new word *w* via a baptism ('This stuff is jade'),

---

<sup>22</sup> The reason for the qualification 'monomorphemic' is that it may be possible for the first use of a morphologically-complex form to be metaphorical. One might never have heard 'hardball' used other than metaphorically, and one could understand metaphorical uses on the basis of knowing how the words 'hard' and 'ball' are used literally, with some help from context. Indeed, for all we know, the first use of this compound may have been metaphorical. (Thanks to Y and J for pushing us on this point.)

stipulation or a mistake that catches on, as in the ‘Madagascar’ example.<sup>23</sup>

Continuation of the originating use requires the intention to use as word *w* with the semantic properties, whatever they are, the originators invested it with. A chain of uses of a word which starts with an originating use and contains only continuing uses is a conservative chain. Conservative uses of *w* are intended to preserve the semantic properties of the originating use. This feature of conservative uses is, we submit, the core of making literal utterances. A name is, *if its use is part of conservative chain*, a ‘channel for the acquisition of knowledge’ (Williamson 2007: 264) about an object and this makes its use literal and scientifically important. Similar things hold for other kinds of expression.

Sainsbury (2005: 117) argues that a continuing use of a proper name also needs to be sensitive to information from the initiating use: some information associated with the name at the point of its introduction is ‘carried forward’. However, there is not a privileged piece of information that needs to be carried forward in every use that continues a practice. Sainsbury’s condition seems plausible for other kind of expressions. There is information that is more firmly associated with a word like the general term ‘cat’: the belief that all cats are animals is firmly held and important for our conception of cats, but we can envisage it being falsified. We can envisage finding out that what we call ‘cat’ are robots controlled from Mars.<sup>24</sup> In this situation we could say ‘Hey, cats turned out not be animals at all!’. What ensured continuation of the ‘cat’ usage was the intention to conform with the originators and preservation of some information (say about shape and movement) associated with the originating use.

Conservative chains of use contrast with *creative* chains:

Creative Chain: the word is originated and it is passed on from speaker to speaker; but some uses are *not* made with the intention to conform to the people from which she acquired the word, yet these uses depend (and are intended to depend) on the originating uses and are sensitive to some information associated with such uses.

---

<sup>23</sup> See Sainsbury 2005: ch. 3.5 for an instructive model of originating uses.

<sup>24</sup> Putnam 1962: 660f.

Take an example. The English noun ‘the sun’ originated as a name of, well, the sun. We can easily use it with the intention to conform to the people who originated it. If we do so, the proper name is part of a channel for the acquisition of knowledge about the sun. But when, for example, Shakespeare wrote:

(12) It is the east, and Juliet is the sun

He neither intends to conform to the original use of ‘the sun’ – he does not want to use it as the originators used it – nor would it be charitable to impute such an intention to him. He would be massively mistaken if he thought that Juliet is the same thing as the star at the centre of the solar system.

But he is also not intending to create a new word by using ‘the sun’ in a new way as someone who dubs his child with the name ‘David’ does. Shakespeare used ‘the sun’ *because* that word already had an established use in which it refers to the sun. Without such an established use, non-literal utterances are not possible. This is the plausible core of the ‘folk concept’ of the non-literal to which Recanati drew our attention. So on the one hand Shakespeare’s use of ‘the sun’ depends on the original use of ‘the sun’ as a term for the sun, while, on the other hand, Shakespeare does not intend to conform to this use.

In what way does Shakespeare’s non-conformist use of ‘the sun’ depend on the originating use? Here we agree with much of the relevance-theoretic view set out above. The original use and conforming uses of ‘the sun’ are associated with a body of beliefs about the referent of the name. We can think of them as collected into a file labelled ‘the sun’. Shakespeare’s use of ‘the sun’ is sensitive to some information in the original file.<sup>25</sup> For example, the original file will contain features like STRIKING, CENTRAL. If – in the context of use – Juliet is neither striking nor central, yet Shakespeare continues to assert that she is the sun, we might start to suspect that he is coining a new word. This suspicion turns out to be unwarranted, if his use of ‘the sun’, though non-conformist, is subject to some constraints that are associated with

---

<sup>25</sup> Note that our claim about sensitivity is compatible with metaphorical uses expressing ‘emergent’ features that are possessed by neither the metaphor vehicle nor the metaphor topic in isolation from each other. On this issue see Wilson & Carston 2006.

the originating uses. The tension between non-conformist intention of the speaker on the one hand and intention to respect some information from the originating uses needs to be resolved. The speaker wants to riff on an existing practice, not start a new one. We need to ascribe to the use of the word a meaning that respects the non-conformist speaker intention and the fact that the speaker did not intend to coin a new word.

With this in mind, we can formulate the *Non-Conformity View of Non-Literal Use* ('NCV' for short) as follows. The basic idea is that a literal use of a word  $w$  is a use of it that is part of a conservative chain; a non-literal use, in turn, is a use of  $w$  that neither an originating use of  $w$  nor made with the intention to conform to such a use, yet the use is sensitive to some information associated with the originating use of  $w$ .

Let us consider an example to illustrate *NCV*. Consider (4) again:

(4 repeated) The mind is a computer.

Suppose the speaker of (4) intends to conform to the original usage of 'computer'. In this case the speaker 'opened up an information channel' to computers. Hence, accepting the utterance of (4) requires us to revise something we may have taken for granted, namely that only artefacts can be computers. Then the utterance may be literally true and if so may tell us something new about computers. Here we have a conservative modulation of 'computer'. Some of the users may be mistaken about what it takes to conform to such a use. However, as long as all uses are made with the intention to conform to the original uses, the name is used literally and utterances containing it can be assessed for literal truth and falsity.

If, on the other hand, (4) is used *without the intention to conform* to the originating use of 'computer', the knowledge channel to computers is blocked. Someone who, for instance, fully understands Shakespeare's utterance of 'Juliet is the sun' cannot thereby extend their knowledge about the sun. This gets at the feature that diminished the value of metaphorical assertions in science. Recall Fodor's (2007) point: "That Juliet is the sun is, perhaps, figuratively true; but since it is only figuratively true, it's of no astronomical interest." On our view, for example, 'The brightest object visible from earth is the sun' is about the sun, while 'Juliet is the sun' is indeed not about the sun and so not a contribution to astronomy.

While live metaphors do, dead metaphors don't meet condition (iii) above. A dead metaphor came into being at some point because there was an originating use, but now it is no longer sensitive to the information associated with the originating use. Understanding a use of 'eye of the needle' does not depend on bringing information about eyes etc. to bear.<sup>26</sup> This is a good result, we think: dead metaphors are as little metaphors as toy ducks are ducks.<sup>27</sup>

## **6. Advantages of the Non-Conformity View of Non-Literal Use**

Let's first make clear why NCV is worth exploring and defending. We agree that the majority of uses of a word express a meaning that is distinct from its standing meaning (if it even has one). Yet we don't take that to be a reason to give up the literal/non-literal distinction. Why? Many of the meanings that are generated 'on the fly' are expressed by uses of the word in which the speaker intends to conform to the originating use. We often need to stretch the use of word to say something scientifically valuable about a topic. In these uses we still want to say something about the original extension of the word, yet we want our audience to correct for contingent and distracting associations which might narrow or widen the extension. These uses are literal, even if they are distinct in meaning from the originating use.

To sum this up: The basic selling point of the Non-Conformity View of Non-Literal Use is that it makes room for lexical modulation, polysemy etc. in literal uses. That is, we have recast the traditional view of metaphor as usage that is in some sense deviant while making room for the (apparent) facts that literal meanings of a word are multiple, contextually sensitive etc.

The Non-Conformity View of Non-Literal Use has further attractive features. It applies neatly to cases that are problematic for the famous Gricean criterion. Grice saw metaphors as violations of a maxim of truthfulness. On this view, the blatant falsehood of the literal meaning of the sentence licences a metaphorical interpretation. One well-known problem for this view is the existence of 'twice-apt' cases like (13), which could be intended (and understood) as metaphorical, even though the sentence has a true literal meaning:

---

<sup>26</sup> The point and this example are from Rey 2017: §4.3.

<sup>27</sup> We agree on this point with Black 1977: 439.

(13) Moscow is a cold city

Our criterion fares better. If an utterance of (13) is intended metaphorically then it is not about coldness. As Fodor might have put it, it's not a contribution to meteorology. Contrast with a literal use, where you could indeed learn that Moscow is cold, or, if you already know about temperatures in Moscow, about how the word 'cold' is used literally.

Many negated metaphors are similarly problematic for the Gricean view, but they are covered unproblematically by our criterion. Consider (14):

(14) No man is an island.

If this is intended metaphorically, it's not about islands: it's not a contribution to geography. Contrast with a similar sentence used literally, which is about islands and from which you *could* learn about geography:

(15) The Isle of Dogs is not an island.

Thus far we have focussed on metaphor, but a second attractive feature of NCV is that it categorises metonymy as non-literal without any need for added stipulations.

Consider some classic cases. The first use of '10 Downing Street' to refer to a group of people who work at 10 Downing Street relied on the already extant tradition in which the name refers to a building, but wasn't intended to belong to that tradition.<sup>28</sup> Similarly with 'The ham sandwich is parked out back' (Nunberg 1995): you can't learn about ham sandwiches from this use (except indirectly, by learning about people who consume them). It would be a mistake to infer from it that ham sandwiches drive cars. The information channel is blocked, just as it is in metaphorical uses.

NCV also has the theoretical virtue that it does not make any assumptions about a metric of 'departure' from the encoded meaning to determine whether a use is non-literal (Wilson and Carston's suggestion). We don't know how to measure the extent of departure and we don't need to: non-literalness comes with lack of intention

---

<sup>28</sup> This particular metonymic use has become conventionalised, but not dead. It is a distinct tradition that still maintains a relation to the literal one.

to conform to originating use. At the same time we can accept with Wilson and Carston that the interpretative mechanism for arriving at the meaning of literal and non-literal uses is the same. But that does not imply that there is no distinction between literal and non-literal uses. The same mechanism is harnessed to satisfy different intentions.

We also don't need to make any substantial assumptions about the standing meaning of a word. For example, we don't need to implausibly assume that there is such a thing in the first place or that it is determined by convention. All we need is the less controversial thesis that there are originating uses and intentions to conform or lack thereof. The originating use may even be empty and yet subsequent uses may be made with the intention to conform to it.

Is the rescued literal/non-literal distinction the one common sense makes? Maybe it is not exactly the same distinction, but a close successor. We aim to articulate the distinction that is of importance in science and semantics and which is tied to the preservation of subject-matter in scientific discourse. This distinction is worth preserving and should not join other babies that are thrown out with the bathwater.

## **7. Denominal Verbs**

We conclude by sharpening our proposal in the light of a *prima facie* counter-example. So-called 'denominal' verbs are derived from nouns. Consider an often used example:

(16) The boy porched the newspaper. (Clark & Clark 1979)

Is this a derivative use of the word 'porch'? If it is, it is a use in a creative chain. The use of the word depends on an originating use, but it is clearly part of a creative chain: the user does not want to go on as before. So it may seem that NCV classifies such uses as non-literal. Yet is 'The boy porched the newspaper' a metaphor? Our intuitions at least don't pronounce clearly on this case.

The right response to this objection seems to us to deny the assumption that the noun 'porch' and the verb 'to porch' are the same word. After all, NCV is only concerned with literal and non-literal uses of *one and the same word*. The verb 'to porch' is a new word whose introduction was possible because 'porch' (noun) was



already around and its introduction is inspired etc. by it. A subsequent utterance of ‘I porched the newspaper’ can be a normal literal use because ‘to porch’ is used with the intention to conform to the initiating uses of ‘to porch’.

Why is the verb a new word although it could only be introduced and used in virtue of the fact that the noun ‘porch’ was already established? The verb has a semantic value of a different kind and different distributional properties. Intuitively, ‘to porch’ was also introduced with the aim of filling a gap in our vocabulary. There is a certain activity related to porches. The innovating speaker creates an expression that signifies this activity by making a new verb, but one that bears on its sleeve its relation to the noun. Hence, denominal verbs don’t threaten our proposal.

There will be other challenges. But we hope to have given the reader good reasons to seek answers to them.

## References

- Belleri, D. (2020). Slurs: Departures from genuine uses and derogation. *Studies in Logic, Grammar and Rhetoric*, 62(1), 9–24.
- Black, M. (1977). More about metaphor. *Dialectica*, 431–457.
- Camp, E. (2006). Metaphor and that certain ‘je ne sais quoi’. *Philosophical Studies*, 129(1), 1–25.
- Carnap, R. (1947). *Meaning and Necessity: A Study in Semantics and Modal Logic*. Chicago: University of Chicago Press.
- Carston, R. (1997). Enrichment and loosening: Complementary processes in deriving the proposition expressed. *Linguistische Berichte*, 8, 103–127.
- Carston, R. (2010). XIII-Metaphor: Ad hoc concepts, literal meaning and mental images. *Proceedings of the Aristotelian Society*, 110(3), 295–321.
- Carston, R. & Powell, G. (2006). Relevance theory: New directions and developments. In E. LePore & B. C. Smith (Eds.), *The Oxford Handbook of the Philosophy of Language* (pp. 341–360). Oxford: Oxford University Press.
- Carston, R. & Wearing, C. (2011). Metaphor, hyperbole and simile: A pragmatic approach. *Language and Cognition*, 3(02), 283–312.
- Chalmers, D. (2008). Foreword. In Clark, A. *Supersizing the Mind: Embodiment, Action, and Cognitive Extension* (pp. ix–xvi). Oxford: Oxford University Press.
- Chomsky, N. (1992). Explaining language use. *Philosophical Topics*, 20, 205–231.
- Clark, E. V. & Clark, H. H. (1979). When nouns surface as verbs. *Language*, 55(4), 767–811.
- Collins, J. (2017). The copredication argument. *Inquiry*, 60(7), 675–702.
- Colyvan, M. (2010). There is no easy road to nominalism. *Mind*, 119(474), 285–306.
- Epstein, R. (2016). The empty brain. *Aeon*, May. <https://aeon.co/essays/your-brain-does-not-process-information-and-it-is-not-a-computer>

- Evans, G. (1973). The causal theory of names. *Proceedings of the Aristotelian Society, Supplementary Volumes* 47, 187–225.
- Fodor, J. A. (1981). The present status of the innateness controversy. In *Representations: Philosophical Essays on the Foundations of Cognitive Science* (pp. 257–316). Cambridge, MA: MIT Press.
- Fodor, J. A. (1987). *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*. Cambridge, MA: MIT Press.
- Fodor, J. A. (1998). *Concepts: Where Cognitive Science Went Wrong*. Oxford: Clarendon Press.
- Fodor, J. A. (2007). Where is my mind? Review of *Supersizing the Mind: Embodiment, Action and Cognitive Extension* by Andy Clark. *London Review of Books*, 29(10), 9–10.
- Fodor, J. A., Garrett, M. F., Walker, E. C. T. & Parkes, C. H. (1980). Against definitions. *Cognition*, 8(3), 263–367.
- Frisson, S. (2009). Semantic underspecification in language processing. *Language and Linguistics Compass*, 3(1), 111–127.
- Frisson, S. (2015). About bound and scary books: The processing of *book* polysemies. *Lingua*, 157, 17–35.
- Giora, R. (2008). Is metaphor unique? In *The Cambridge Handbook of Metaphor and Thought* (pp. 143–160). Cambridge: Cambridge University Press.
- Glanzberg, M. (2008). Metaphor and lexical semantics. *The Baltic International Yearbook of Cognition, Logic and Communication*, 3, 1–47.
- Gross, S. (2021). Linguistic judgments as evidence. In N. Allott, T. Lohndal & G. Rey (Eds.), *A Companion to Chomsky* (pp. 544–556). Hoboken, NJ: Wiley Blackwell.
- Harris, D. W. (2020). Semantics without semantic content. *Mind & Language, Early view*.
- Jeshion, R. (2013). Expressivism and the offensiveness of slurs. *Philosophical Perspectives*, 27(1), 231–259.
- Lande, K. (2019). Do you compute? *Aeon*. <https://aeon.co/essays/your-brain-probably-is-a-computer-whatever-that-means>
- Lewison, K. (2015). Metaphors and legal reasoning. *The Chancery Bar Association Lecture 2015*: <https://www.chba.org.uk/for-members/library/annual-lectures/metaphor-legal-reasoning>
- Lossius Falkum, I. & Vicente, A. (2015). Polysemy: Current perspectives and approaches. *Lingua*, 157, 1–16.
- McNally, L. & Spalek, A. (2019). ‘Figurative’ uses of verbs and grammar. *Unpublished MS*.
- Noveck, I. A. (2018). *Experimental Pragmatics: The Making of a Cognitive Science*. Cambridge: Cambridge University Press.
- Nudds, M. (2004). The significance of the senses. *Proceedings of the Aristotelian Society* 104, 31–51.
- Nunberg, G. (1995). Transfers of meaning. *Journal of Semantics*, 12(2), 109–132.

- Pietroski, P. (2003). Small verbs, complex events: Analyticity without synonymy. In L. M. Antony & N. Hornstein (Eds.), *Chomsky and His Critics* (pp. 179–214). Malden, MA: Blackwell.
- Putnam, H. (1962). It ain't necessarily so. *The Journal of Philosophy*, 59, 658–71.
- Quilty-Dunn, J. (2020). Polysemy and thought: Toward a generative theory of concepts. *Mind & Language*, 1–28.
- Quine, W. V. (1981). *Theories and Things*. Cambridge, MA: Harvard University Press.
- Ramchand, G. (2014). On structural meaning vs conceptual meaning in verb semantics. *Linguistic Analysis*, 39(1), 207–244.
- Recanati, F. (2004). *Literal Meaning*. Cambridge: Cambridge University Press.
- Rey, G. (2017). The analytic/synthetic distinction. *Stanford Encyclopedia of Philosophy*.
- Richard, I. (2014). Metaphors in English for law: Let us keep them. *Lexis [online]*, 8, 1–19.
- Rodd, J. M. (2020). Settling into semantic space: An ambiguity-focused account of word-meaning access. *Perspectives on Psychological Science*, 15(2), 411–427.
- Sainsbury, R. M. (2005). *Reference Without Referents*. Oxford: Clarendon Press.
- Schroeter, L. & Schroeter, F. (2016). Semantic deference vs semantic coordination. *American Philosophical Quarterly*, 53, 193–210.
- Searle, J. R. (1980). The background of meaning. In J. R. Searle, F. Kiefer & M. Bierwisch (Eds.), *Speech Act Theory and Pragmatics* (pp. 221–232). Dordrecht: Springer.
- Searle, J. R. (1983). *Intentionality: An Essay in the Philosophy of Mind*. Cambridge: Cambridge University Press.
- Spalek, A. A. (2015). Spanish change of state verbs in composition with atypical theme arguments: Clarifying the meaning shifts. *Lingua*, 157, 36–53.
- Sperber, D. & Wilson, D. (1986). *Relevance: Communication and Cognition* (2nd ed. 1995). Oxford: Blackwell.
- 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: Cambridge University Press.
- Sperber, D. & Wilson, D. (2002). Pragmatics, modularity and mind-reading. *Mind and Language*, 17(1–2), 3–23.
- Stern, J. (2006). Metaphor, literal, literalism. *Mind & Language*, 21(3), 243–279.
- Tiersma, P. M. (2001). A message in a bottle: Text, autonomy, and statutory interpretation. *Tulane Law Review*, 76, 431–482.
- Vicente, A. (2019). Chomskyan arguments against truth-conditional semantics based on variability and co-predication. *Erkenntnis*, 1–22.
- White, R. M. (1996). *The Structure of Metaphor: The Way the Language of Metaphor Works*. Oxford: Blackwell.
- White, R. (2001) Literal meaning and “figurative meaning”. *Theoria* 67, 24–59.
- Williamson, T. (2007). *The Philosophy of Philosophy*. Oxford: Blackwell.

- Wilson, D. & Carston, R. (2006). Metaphor, relevance and the ‘emergent property’ issue. *Mind & Language*, 21(3), 404-433.
- 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). Basingstoke: Palgrave Macmillan.
- Wilson, D. & Sperber, D. (2002). Truthfulness and relevance. *Mind*, 111(443), 583–632.
- Yablo, S. (1998). Does ontology rest on a mistake? *Aristotelian Society Supplementary Volume* 72(1), 229–261.