Settling into Semantic Space: An Ambiguity-Focused Account of Word-Meaning Access

Jennifer Rodd

Department of Experimental Psychology, University College London

j.rodd@ucl.ac.uk

@jennirodd

www.jennirodd.com

THIS PAPER IS A DRAFT AND HAS NOT BEEN THROUGH PEER REVIEW

30.1.2019

Abstract

Most words are ambiguous: individual wordforms (e.g., “run”) can map onto multiple different interpretations depending on their sentence context (e.g., “the athlete/politician/river runs”). Models of word-meaning access must therefore explain how listeners and readers are able to rapidly settle on a single, contextually appropriate meaning for each word that they encounter. This article presents a new account of word meaning that places semantic disambiguation at its core, and integrates evidence from a wide variety of experimental approaches to explain this key aspect of language comprehension. The model has three key characteristics. (i) Lexical-semantic knowledge is viewed as a high-dimensional space; familiar word meanings correspond to stable states within this lexical-semantic space. (ii) Multiple linguistic and paralinguistic cues can influence the settling process by which the system resolves on one of these familiar meanings. (iii) Learning mechanisms play a vital role in facilitating rapid word-meaning access by shaping and maintaining high quality lexical-semantic knowledge. Several key areas for future research are identified.

Word Meaning Access: The Challenge of Lexical Ambiguity

The ability to rapidly and accurately access the meanings of words when they occur within sentence contexts is a key component of language comprehension. This task is made difficult by the inherent ambiguity of most words, which can refer to different concepts in different contexts. This review integrates current research to present a unified theoretical account of on how ambiguous words are learned, represented and processed.

The most salient form of lexical ambiguity is found in homonyms such as “trunk” that have multiple unrelated meanings (e.g., “the trunk of a car/tree/elephant). This form of ambiguity is relatively rare, and is present for about only 7% of relatively frequent wordforms (Rodd, Gaskell, & Marslen-Wilson, 2002). In contrast, more than 80% of wordforms are polysemous - they can refer to more than one related word sense (Rodd et al., 2002). For example, the verb “run” is highly polysemous – it has a multitude of different interpretations that are appropriate within different sentence contexts (e.g., “the athlete/river/program/paint/manager/dye/train/candidate runs”). Successful word-meaning access occurs when the *correct* interpretation (i.e. the interpretation that was intended by the speaker/writer) is selected from the range of familiar possibilities.

Lexical ambiguity is often viewed within psycholinguistics as a troublesome nuisance that adds to the processing demands that are associated with successful language comprehension (Johnsrude & Rodd, 2016). However it is important to remember that polysemy adds considerably to the communicative power of language by allowing communicators to flexibly convey a rich array of meaning from their finite set of familiar word forms. Polysemy also provides a source of linguistic creativity: speakers can creatively extend familiar words beyond their original meanings. This creative aspect is most apparent for ‘regular polysemy’: those clusters of words that have systematic patterns of senses. For example, the meanings of animal names can be productively extended to refer to the meat that comes from that animal (e.g., chicken, lamb, ostrich etc.; Copestake & Briscoe, 1995; Mahesh Srinivasan & Rabagliati, 2015). In addition, it has been argued that ambiguity is a functional property of language that *improves* communicative efficiency (Piantadosi, Tily, & Gibson, 2012). Specifically, these authors argued, based on a statistical analysis of several different languages, that so long as words are used within rich, informative contexts, then this ‘re-use’ of wordforms actually *reduces* the processing demands on the language system. (Piantadosi et al., 2012).

This acknowledgement that most common words are (to varying degrees) ambiguous, requires that lexical-semantic disambiguation is a core component of models that aim to explain how word meanings are accessed/retrieved. This approach is in stark contrast with many influential models of lexical processing, which make the simplifying assumption that there is a one-to-one mapping from word form to word meaning

(e.g., Plaut, McClelland, Seidenberg, & Patterson, 1996) and provide no mechanism for dealing with lexical-semantic ambiguity.

One highly influential model of word-meaning access, that did explicitly acknowledge the challenges placed by lexical ambiguity on word-meaning access, is the ‘Reordered Access’ model (Duffy, Morris, & Rayner, 1988). The model made two central claims. First, when a wordform is encountered, its different, familiar meanings become available simultaneously. Second, the timing with which specific meanings become available is influenced by two factors: (i) their relative frequency in the language (dominance), and (ii) the preceding sentential context. These constraints allow for rapid access (and then integration) of word meanings that are highly frequent in the language and/or that are strongly supported by the surrounding context. These properties allow the word-meaning access system to avoid unnecessary distraction from word meanings that are unlikely to be correct. These two core claims of the Reordered Access model have stood the test of time and are supported by a wide range of evidence, for both written and spoken language (see Vitello & Rodd, 2015 for review).

More recently however, research has revealed a richer and more complex view of how the meanings of ambiguous words are learned, represented and processed. While these results do not undermine the core claims of the Reordered Access model, evidence has now shown that readers/listeners use a far wider range of distributional cues about the usage of different words meanings and make use of this learned information to ‘nudge’ themselves towards correct meanings quickly and accurately. The current paper will review current experimental findings, and integrate these findings into a new empirically driven theoretical account of word-meaning access that has three key characteristics.

i) Distributed representations of word meanings

Familiar word meanings are represented by distributed representations that correspond to stable states within a complex, structured high-dimensional semantic space (Armstrong & Plaut, 2016; Rodd, Gaskell, & Marslen-Wilson, 2004).

ii) Fluent disambiguation is facilitated by integration of multiple linguistic and paralinguistic cues

A wide range of contextual cues influence word-meaning access, such that contextually appropriate meanings are more readily available. The immediate sentence context provides the primary disambiguating cue but, word-meaning access is situated within a highly interactive cognitive system that also allows non-linguistic cues (such as the identity of the speaker) to support successful word-meaning access (Cai et al., 2017).

iii) Learning mechanisms shape and maintain high quality lexical-semantic knowledge

Lexical-semantic space continues to be shaped by personal linguistic experience throughout the lifespan. Not only are new *unfamiliar* word meanings being continually integrated into the existing lexical-semantic space (e.g., Rodd et al., 2012), but also learning from recent linguistic experiences with *familiar* word meanings continues to reshape and maintain high quality lexical knowledge throughout the lifespan (Rodd, Lopez Cutrin, Kirsch, Millar, & Davis, 2013).

Distributed Representations of Word Meanings

Word meanings are dynamic and highly variable. The same word form can map onto to a multitude of different semantic interpretations. One way of capturing this is to assume that each familiar word meaning corresponds to a single point in a high dimensional semantics space (Armstrong & Plaut, 2008; Borowsky & Masson, 1996; Joordens & Besner, 1994; Kawamoto, Farrar IV, & Kello, 1994; Rodd et al., 2004).

Within this framework, the meaning of a relatively unambiguous word (e.g., “SHOE”) will corresponds to the single point within semantic space in which all its constituent semantic features are active and all other possible semantic features are inactive. The form-to-meaning mapping for such words is straightforward and unambiguous: a single wordform maps directly onto its single lexical-semantic representation (Figure 1). In

contrast, for ambiguous words the situation is more complex: a single wordform maps forward onto multiple different interpretations that consist of different combinations of semantic features. It is the one-to-many mapping between form and meaning that creates the need for semantic disambiguation. This view that single wordforms can map onto many different possible lexical-semantic representations is in direct contrast with a more common computational approach to modelling the mapping from wordform to word meanings that usually make the simplifying assumption that the meanings of ambiguous wordforms can be captured by a single vector within semantic space (e.g., Plaut, McClelland, Seidenberg, & Patterson, 1996).

Importantly, this approach captures a key linguistic distinction between homonymy and polysemy (Rodd et al., 2002). Homonyms (e.g., “bark”) that have multiple, unrelated meanings will map onto multiple uncorrelated combinations of lexical-semantic features (Figure 1). These distinct word meanings correspond to distant points within semantic space. In contrast, polysemous words (e.g., “run”) consists of a single wordform that maps onto multiple different correlated representations that are situated within a single area of semantic space (Figure 1). Thus, within this framework the distinction between related words senses and unrelated word meanings is therefore simply one of degree. In both cases the different interpretations of the wordform correspond to different points within semantic space. The only difference is the relative proximity of the meanings: related words senses share many semantic features and therefore lie in adjacent areas of semantic space while unrelated word meanings are uncorrelated and therefore correspond to more distant locations. This approach accommodates the finding that there is considerable variability in relatedness within the class of polysemous words (see Klein & Murphy, 2001 for discussion of how different senses of polysemous words are often be judged to be relatively *unrelated*).

Figure 1. Representations of different types of words within distributed framework. (Drawings taken from Betts, Gilbert & Rodd, in prep.)

A critical concept within this framework that allows a model of this type to cope with the prevalence of ambiguity is that of the attractor basin. As the system learns the appropriate form-to-meaning connections lexical-semantic space develops a highly structured collection of interconnections such that semantic features that tend to co-occur are positively connected, while incongruent semantic features develop reciprocal inhibitory connections. These connections ensure that familiar combinations of lexical-semantic features that correspond to known word meanings correspond to highly stable states (known as attractor basins). Once the network enters such an attractor basin its activation tends to remain relatively unchanged – there is no pressure for the network to move away from that stable combination of semantic features (until the next word is encountered). In contrast, if the network activates an unfamiliar/inappropriate combinations of semantic

features, the co-activation of incongruent semantic features ensures that its activation state is not stable. Under such conditions the network of connections within the lexical-semantic representations will ensure that the activation of semantic features will update/change in a manner that moves its current state away from a meaningless, unstable state towards a more stable, familiar representation.

This distributed framework views the retrieval of a word's meaning as a dynamic settling process. In particular, for ambiguous words, the lexical semantic network will usually initially enter a relatively unhelpful blend state that will likely contain inconsistent elements of meanings that correspond to different possible interpretations of the incoming wordform. The network then incrementally moves away from this unstable state and settles into a stable state that corresponds to one of the familiar interpretations of that word. Successful word meaning retrieval occurs when the network successfully settles into a familiar, stable state within lexical semantic space. Various factors such as the prior state of the network and other currently active semantic information (i.e. its preceding context) as well as recent/long-term experience with the ambiguous word itself will determine which of the alternative meanings is more robustly activated and wins the competition during the active settling process. This attractor structure therefore provides the core mechanism for semantic disambiguation, ensuring that whenever an ambiguous word is encountered one or other of its possible interpretations is eventually settled upon and it is unlikely to linger in a non-meaningful blend state.

A model of this sort was previously implemented by Rodd et al. (2004) building on earlier connectionist models of lexical ambiguity (Borowsky & Masson, 1996; Joordens & Besner, 1994; Kawamoto, Farrar & Kello, 1994). Simulations using this network demonstrated that the network showed appropriate settling behaviour, which accurately simulated human lexical decision performance: settling times were delayed for words that had multiple different unrelated meanings due to the additional processing required to move the network from an initial blend state that corresponded to a mixture of the words' two meanings towards one or other settled familiar word interpretation. For polysemous words, a more complex picture emerged: early in the settling process, these words benefitted from their broad, deep semantic attractor basis that corresponded to the large cluster of semantically unrelated meanings. However later in settling, a disadvantage emerged as the ambiguity between the different related senses prevented optimal settling behaviour (See Armstrong & Plaut, 2016 for extensive discussion).

Future work is clearly needed to develop the scope of these computational/cognitive models. In particular, work is needed to reveal how the meanings of multiple consecutive words can be simultaneously represented within the same network, most likely by allowing the semantic features of each word to remain separately bound to its corresponding phonological/orthographic features (Rabagliati, Dumas, & Bemis, 2017).

Fluent disambiguation is facilitated by fluent integration of linguistic and paralinguistic cues

Within this theoretical framework, meaning disambiguation is viewed as a dynamic settling process in which the current state of the network moves away from some initial 'blend state' that includes elements of multiple different interpretations, towards a single familiar word meaning. However, while the presence of attractor structure within lexical-semantic knowledge ensures that this settling process always occurs, on its own this property of the network cannot ensure that the 'correct' meaning is accessed. Fluent, accurate semantic disambiguation requires that the network is highly sensitive to the distributional statistics of language, such that during the settling process the network's activation is 'nudged' away from irrelevant meanings towards the correct interpretation of the wordform: the network must utilise a variety of predictive cues to settle into the correct attractor basin and not be lured into those basins that correspond to irrelevant distractor meanings.

The most salient cue that helps readers/listeners to access the correct meaning of each wordform is the presence of biasing words in the preceding sentence context. For example, the word "bark" will be interpreted very differently when it is preceded by either the word "dog" or the word "tree". A large body of research has attempted to specify exactly *how* these cues operate. The field has focused on three key (related) questions about how sentence context influences word meaning access. Specifically, researchers have attempted to specify (i) the stage of processing at which context plays a role, (ii) the mechanism(s) by which the system

assesses the congruency between a word's alternative meanings and the sentence context, and (iii) the type(s) of contextual information that can influence word-meaning access. The following sections reviews the relevant evidence on these key issues and how these findings can best be accommodated within the current distributed framework.

The psycholinguistic literature contains a wealth of evidence from studies that aimed to reveal the stage of processing at which contextual cues operate (see Witzel & Forster, 2015). These experiments have aimed to distinguish between opposing views about the structural relationship between (low-level) word-meaning access and (higher-level) sentence processing mechanisms. At one end of the theoretical spectrum are '*exhaustive access*' models of word meaning access (Onifer & Swinney, 1981; Seidenberg, Tanenhaus, Leiman, & Bienkowski, 1982; Swinney, 1979). These models fall within the general class of *autonomous* or *modular* models of language processing (Fodor, 1983). Under this view, when an ambiguous wordform is first encountered, all of its familiar word meanings are all automatically initially accessed/retrieved in parallel regardless of any contextual cues. Such models only permit context to influence the meaning selection/integration processes that occur after all the possible meanings have been accessed/retrieved. In contrast, '*selective access*' models fall within the class of *interactive* models of word-meaning access, and assume that contextual cues can directly modulate word-meaning access. Under the strongest version of this interactive account, if the preceding context is sufficiently strong then it can prevent the access/retrieval of meanings that are inconsistent with this context and can thereby ensure that only the contextually relevant meaning is accessed (Tabossi, 1988; Tabossi, Colombo, & Job, 1987).

Taken together, the available evidence is incompatible with the most extreme versions of both these accounts (Sheridan, Reingold, & Daneman, 2009). First, the modular position that context can make *no* contribution to the initial access of word meanings seems incompatible with a range of findings showing that the time taken to read ambiguous words varies considerably according to whether the preceding context is supportive or inconsistent with participant word meanings. (See Duffy, Kambe, & Rayner, 2001 for a comprehensive review of findings concerning how ambiguous words are processed within sentence contexts.) In addition, the strongest version of the interactive models are inconsistent with the finding (known as the 'subordinate bias effect') that readers show consistent processing delays when an ambiguous word with one strongly dominant meaning (e.g., "pen") are preceded by words that are strongly biased the reader's interpretation towards a different, subordinate meaning (e.g., "Because it was too small to hold all the new animals, the old pen was replaced"; Duffy et al., 1988). This effect is usually interpreted as showing that prior context *cannot* prevent the reader from automatically accessing the most frequent word meaning (e.g., Kambe, Rayner, & Duffy, 2001; Kellas & Vu, 1999; Pacht & Rayner, 1993; Rayner, Binder, & Duffy, 1999; Rayner, Pacht, & Duffy, 1994; Sereno, 1995).

Researchers have therefore converged on an intermediate view, as instantiated in the reordered access model: all familiar word meanings are initially accessed, but context can directly influence their relative levels of activation/availability such that word meanings that are compatible with the preceding sentence context are more readily available than they would be in an unsupportive context (Duffy et al., 2001). In the context of the current distributed framework, this view is operationalised by assuming that contextual cues can directly modulate the settling behaviour of the network. Specifically, when an ambiguous wordform is encountered activation feeds forward to automatically begin to activate semantic features that correspond to all of its familiar interpretation: even strong context cannot prevent the network from initially entering a 'blend state' that contains some inappropriate meaning elements. Under this view the primary role of context is to 'nudge' the current state of the network during the setting process such that the network moves rapidly towards the word meaning that is supported by the current context. This approach allows context to influence (i) which meaning is ultimately selected (settled into), as well as (ii) how rapidly this settling process occurs. Importantly, this approach does *not* rule out the possibility that when relatively balanced ambiguous words occur within a neutral context, then settling may be sufficiently slow that words that occur *after* the ambiguity can influence the settling trajectory. For example, in the spoken sentence "The woman hoped that both pears/pairs taste sweet", the initial settling process for the homophone "pears/pairs" may be influenced by the meaning of the word "taste" (Rodd, Longe, Randall, & Tyler, 2010). Current evidence is equivocal about how long the system can remain in this transient pre-disambiguation 'blend' state, i.e. how long it can

maintain elements of meaning that correspond to more than one possible interpretation (Vitello, 2013). Resolving this issue remains a key challenge for the field.

One key property of this proposed mechanism by which context can directly influence word-meaning access is that it avoids the need for the system to incorporate a mechanism that directly assesses the relative congruence of all the different meanings of a word within a given context. This approach contrasts with the (computationally demanding) alternative that have been suggested in which the system initially accesses all possible word meanings and attempts (in parallel) to (i) integrate these different possible interpretations with the ongoing representation of the preceding sentence/discourse, and then (ii) compare these (partially) complete higher-level representations to determine which of the alternative word meanings is most plausible. However given that most words are (to varying degrees) ambiguous and a single sentence can therefore contain numerous consecutive ambiguities, such a mechanism could result in a requirement to compute infeasible numbers of complex discourse representations in parallel (see Witzel & Forster, 2015 for further discussion). The current view is fully consistent with the more parsimonious approach put forward by Witzel & Forster (2015), who propose that all meanings of an ambiguous word are initially activated, with subsequent meaning selection being driven by a “fast-acting, low-level heuristic based on intralexical connections” (Witzel & Forster, 2015; pg 159).

This approach that if any of the words that precede an ambiguous word are either (i) strongly associated or (ii) overlapping in semantic content with one of the word’s possible meanings then these contextual cues will act to ‘nudge’ the activation state towards that word meaning. For example, in the phrase “the dog’s bark” the lexical-semantic knowledge activated in response to “dog” will be rapidly retrieved and then will influence the settling trajectory for the word “bark” by nudging it towards the dog-related meaning, due to the learned associative/semantic relationships between the words “dog” and “bark”. Under this view, contextually appropriate meanings benefit from an accumulation of positive reinforcement from all the related semantic content that is currently co-active within the language system, while inappropriate meanings are discarded because the ‘winner-takes-all’ nature of the attractor basin structure necessarily requires them to be deactivated in the presence of a contextually supported alternative. Thus the rejection of inappropriate meanings is viewed as a relatively passive process in which they are essentially discarded due to the absence of evidence in their favour, not because of a more active process that first activates and then discards them. This approach does not require the system to explicitly evaluate the potential relationships between the context and *all* possible meanings.

In natural language environments, words are usually presented within rich conceptual contexts that provide multiple convergent associative/semantic cues, thus it is likely that such a process of low-level associative modulation of lexical-semantic activation is likely to result in the activation of the ‘correct’ word meaning for the vast majority of words that are encountered. However, on some, relatively rare, occasions the system will produce an inappropriate response to a wordform and settle into the wrong attractor basin. On these occasions the ‘wrong’ meaning will be ‘passed forward’ to the subsequent higher-level comprehension processes, i.e. will begin to be integrated with other elements of the sentence/discourse. These (rare) conditions will result in the need for a re-evaluation of the word’s meaning. We will return to this reinterpretation aspect of semantic-disambiguation later.

A final important question concerning how context influences word-meaning access, is the requirement to specify exactly what forms of contextual cues are permitted to influence the settling process. The current, interactive account takes a pragmatic approach: it assumes that *any* conceptual/semantic information that is currently active within the language system (broadly construed) can potentially influence lexical-semantic disambiguation, and that the degree of influence is directly related to the current level of activation of the relevant cue. For instance when the ambiguous word “bark” is immediately preceded by the word “dog” the ‘nudge’ provided will be very strong due to (i) the strength of the semantic-associative relationship and (ii) the proximity of the words resulting in an overwhelmingly strong contextual bias that would drive very rapid and successful disambiguation. In contrast, if the only cue to the disambiguation of “bark” was a word that was more weakly represented in the ongoing discourse or was more weakly related/associated to the appropriate meaning then the degree of contextual constraint would be weaker making disambiguation (via settling) a slower and less reliable process.

Importantly, recent evidence has revealed that the meanings of words in the immediate sentence context is not the only source of contextual bias that influences word-meaning access. Recent results from Cai et al., (2017) demonstrate that listeners can keep track of how ambiguous words are used differently in specific linguistic environments. Cai et al., (2017) examined the effect of speaker accent (British versus American) on the interpretation of words that have different dominant meanings in the two dialects (e.g., “bonnet”). Results from a range of tasks confirmed that for native speakers of British English the meaning that is dominant in British English (i.e. car part) is (of course) usually more readily available. But intriguingly, Cai et al., (2017) found that the availability of the subordinate American-dominant meaning (e.g. type of hat) can be boosted when the words are spoken in an American accent. Further, this shift in preference toward American-dominant meanings did not transfer to written words that were intermixed with the presentation of US accented spoken words. These data make clear that listeners are able to keep track of reliable differences in how words are used by speakers with specific accents compared with (unaccented) written English. Importantly, these findings cannot be explained via a simple semantic-priming type mechanism: it is not that all ‘hat related’ meanings are more active when a US accent is detected, it is only in the presence of the word form “bonnet” that the hat-related meaning is boosted in the presence of this accent. In addition, Cai et al., (2017) suggest that speaker accent is unlikely to be ‘special’ in terms of its ability to directly modulate word-meaning availability and that listeners are likely to have also learned expectations about how word-meaning use varies according to other group-level factors such as age, gender and social class. Perhaps most intriguing is the possibility that we may keep track of the idiosyncratic biases in the ways that highly familiar individuals tend to use specific words. Future work is needed to determine the extent to which listeners (and readers) acquire and maintain specific word-meaning usage knowledge of this sort. Further research is also needed to assess the extent to which broader contextual cues (e.g., visual scene information) can directly modulate the initial stages of word-meaning access.

Importance of Learning Mechanisms in the maintenance of high quality lexical representations

Within this framework developing and maintaining high quality lexical knowledge with a sophisticated attractor structure that can appropriately guide semantic disambiguation is critical for fluent word meaning access. Therefore a foundational assumption of this theoretical approach is that learning plays a key role in supporting skilled comprehension of words.

First, the system must (of course) be able to learn entirely new form-meaning pairs that are encountered (e.g., learning for the first time that “ping” refers to getting in touch via digital means). Second, and more relevant to the topic of lexical-semantic disambiguation, the system must retain the capacity to learn new meanings for already familiar wordforms, for example due to technological innovation (e.g., the social media senses of “tweet”, “post” and “troll”). New meanings must also be learned when we learn new academic subjects or hobbies (e.g., the rowing-related meanings of “square”, “catch” and “feather”), or enter a new dialect community (e.g., “piece” can mean a sandwich in Scotland and a gun in America). This form of word learning can result in a previously unambiguous word with a previously straightforward form-to-meaning mapping gaining additional complexity with respect to its lexical-semantic representation and acquiring a more complex one-to-many mapping. Similarly, already ambiguous words can continue to acquire additional meanings/senses throughout their lifespan. Thus, exposure to new word meanings can result in lexical-semantic space being reshaped and gaining complexity. Studies have shown that adult participants are highly skilled at working out the meanings of novel word senses (Clark & Gerrig, 1983; Frisson & Pickering, 2007) and that they can retain good memory for new word meanings after relatively few exposures, with relatively little forgetting being observed after a one week delay (Hulme, Barsky, & Rodd, 2018).

Recent studies of adult word learning have revealed several interesting characteristics of this form of new meaning learning. Fang, Perfetti and Stafura (2017) compared learning of new meanings for familiar words to a condition in which both the form and meaning of the new word were unfamiliar. Their results support a ‘biphasic pattern’ in which early in learning is more successful for familiar wordforms compared to the novel wordforms, but later in learning this pattern is reversed. They suggest that the familiarity of a wordform initially facilitates learning because the relevant form-based representation already exists and so does not need

to be newly constructed, but that later in learning the interference between the alternative meanings plays a more influential role in constraining learning success. Similarly, Rodd et al., (2012) showed that new meanings for familiar wordforms are more easily learnt when they are semantically *related* to a word's existing meaning. For example, it is relatively easy to learn that a "sip" is 'a small amount of data' due to the semantic link with its existing 'small drink' meaning. Taken together these studies indicate that, at least for adults, the pre-existing structure of lexical-semantic space, and the existing form-to-meaning mapping(s) for a particular word form has a significant impact on the ease with which new lexical knowledge can be inserted into the store of lexical-semantic knowledge. All new learning occurs within an already highly structured system – this existing structure is of significant benefit when the to-be-learned information shares informational content with existing knowledge but creates additional challenges when form-meaning pathways must be built that conflict with the information already stored.

In addition to this form of new meaning learning in which an entirely new stable state is inserted into the lexicon, it has recently become clear that learning mechanisms also play a key role in allowing individuals to maintain and reshape their existing lexical semantic representations for *familiar* word meanings throughout their lifespan.

It has long been known that individuals are sensitive to the distributional properties of word meanings. In particular, an extensive body of work has emphasized that both listeners and readers are highly sensitive to the overall relative frequencies of word meanings in the language, known as 'meaning dominance' (e.g., Hogaboam & Perfetti, 1975; see Duffy et al., 2001 for a review). There is abundant evidence confirming that more frequent word meanings are more easily accessible. For example when presented without any surrounding sentence context a word like "pen" will most likely be interpreted as referring to its most common 'writing implement' meanings and not as referring to its less frequent (subordinate) 'animal enclosure' meaning (Twilley, Dixon, Taylor, & Clark, 1994). In addition, as discussed earlier, a large number of studies in which readers' eye-movements are tracked during reading have shown that meaning dominance plays a key role in determining how easily ambiguous words are understood within a range of different types of sentence structures (Duffy et al., 2001; Vitello & Rodd, 2015). These dominance effects can be readily explained within the attractor basin framework by assuming that the mappings from form to meaning are built up over time such that stronger form-to-meaning connections are built up for word meanings that have been more frequently encountered (Mukhia, 2018).

More recently however, a more nuanced picture has emerged in which the availability of word meanings is driven by an individual's specific pattern of *recent* and *longer-term* experiences with that particular word (Rodd et al., 2016). Research has shown that adult lexical-semantic representations are surprisingly malleable and that learning mechanisms play a key role in allowing individuals to keep track of both short-term and long-term changes in their linguistic environment.

First, it has been shown that individuals keep track of *long-term* changes by adjusting their lexical-semantic representations in response to their recent linguistic environment such that those meanings that are more likely to be encountered in the future are most easily accessible. Rodd et al. (2016) investigated word meaning preferences within a group of recreational rowers who all knew a set of rowing-related meanings for highly familiar words such as "catch", "square" and "feather". They found that the availability of these specific meanings increased with the number years of rowing experience of the participants, showing that the availability of these meanings increased in response to individuals' prolonged exposure to an environment in which these word meanings were often used. Similarly, Wiley, George & Rayner (2016) showed that baseball experts had more trouble ignoring word meanings that relate to baseball (compared with non-experts) even within sentence contexts that are strongly biased towards the non-baseball meaning (e.g. "Monica had a great fear of things flying around her head. She looked for the BATS that lived in the shed"). These studies support the view that individuals respond to their long-term experience with particular vocabulary items over periods of several years by increasing the availability of these word meanings. This work also emphasises the individual differences that are inherent in lexical-semantic representations – no two people share the same exact linguistic history and so each individual has their own uniquely idiosyncratic store of lexical-semantic knowledge. While speakers of a shared language will (mostly) share a common set of known word meanings,

the ease with which these different word meanings come to mind will differ significantly according to their specific experiences with these word meanings.

In addition to these incremental effects of long-term experience with word meanings, evidence from word-meaning priming paradigms has revealed listeners and readers are particularly sensitive to very recent encounters with ambiguous words (Rodd et al., 2013): recently encountered word meanings are more readily available than would be expected from their overall frequency/dominance. In particular, lower frequency meanings that are usually relatively difficult to access become more readily available when they have been encountered within the preceding minutes/hours. This ‘boosting’ of recently encountered meanings serves to reduce the processing load when word meanings are used repeatedly within the same conversation or document.

In the standard word-meaning priming paradigm, participants initially encounter ambiguous words within strongly disambiguating contexts (e.g., “The farmer moved the sheep into the PEN”). Then, after a delay (typically 20-40 minutes) they perform a task that provides a measure of the availability of the word’s meanings. The most commonly used test of word-meaning availability has been word association (Betts, Gilbert, Cai, Okedara, & Rodd, 2018; Gilbert, Davis, Gaskell, & Rodd, 2018; Rodd et al., 2016, 2013); in these experiments, for each ambiguous word participants are asked to produce the first related word that comes to their mind (e.g., “PEN-enclosure” vs “PEN-paper”). Results show that participants are more likely to retrieve the subordinate word meaning if it had been previously encountered (compared to a control, unprimed condition). Speeded semantic related task has also been used to measure meaning availability: participants decide whether the ambiguous wordform is related to a probe word that is only related to one of its meanings (e.g., “PEN-enclosure”). Faster responses to the probes related to the primed meaning (compared with an unprimed baseline) are taken as evidence that the primed meaning has become more readily available (Gilbert et al., 2018).

Both these word-meaning priming paradigms have consistently shown that following just one encounter with an ambiguous word in a strongly disambiguating context there is a substantial boost in meaning availability. This boost is particularly strong at very short delays (2-3 Minutes; Rodd et al., 2013; 2016), but remains robust and relatively stable after 20-40 minutes (Rodd et al., 2016). Remarkably, significant priming from a single prime encounter has even been reported to be present after 24 hours when participants are able to sleep soon after the encounter (Gaskell, Cairney, & Rodd, 2019), suggesting that these short-term changes to word meaning availability can make substantial contributions to longer-term changes. Word-meaning priming has not only been shown in carefully controlled lab-based conditions, but also in more naturalistic conditions where the primes were presented within short vignettes within a radio programme and the test of priming was conducted via a subsequent web-based word association task (Rodd et al., 2016).

Word-meaning priming experiments make clear that adults adjust their lexical-semantic knowledge in response to their linguistic environment. They also have revealed several important characteristics of the underlying mechanism(s) that supports this learning. Gilbert et al., (2018) examined the extent to which information about word meanings that is gained within one modality (e.g., speech) is transferred to the other modality (e.g., reading). In these experiments ambiguous target words were primed with either written or spoken sentence. After a 20 minute delay, participants’ interpretation of these ambiguous words was tested by presenting them in either written or spoken form. Experiments using both word association and speeded semantic relatedness decisions as measures of word-meaning priming revealed the same pattern: (i) significant cross-modal priming and (ii) no evidence that cross-modal priming was reduced compared to the condition in which both prime and test were presented in the same modality. These results were interpreted as evidence that priming should be characterised as changes to an amodal store of lexical-semantic knowledge that is shared between listening and reading. Importantly, these data are inconsistent with early accounts of word-meaning priming that suggested that learning resulted in (modality-specific) changes to the mapping between a word’s phonological/orthographic form and its (amodal) semantic representation (Rodd et al., 2013).

However, although these results suggest that learning in the standard word-meaning priming paradigm is amodal in nature, they do not require that all learning about the likely future availability of word meanings is generic and will necessarily transfer to all future instances of that word, regardless of factors such as its

modality. Indeed the effects of speaker accent described earlier (Cai et al., 2017) make clear that listeners can (under some conditions) to keep track of reliable differences in how words are used differently in different forms of spoken English compared with written English. Taken together these two sets of experiments suggest that individuals usually use their linguistic experiences to update modality-general expectations about the likelihoods of different meanings, unless there are salient, systematic cues (such as accent) that indicate that words are being used differently in different environments.

One intriguing question that arises is how learners make changes to their lexical-semantic knowledge into the lexicon *without* negative consequence for their existing stored knowledge. One explanation of how new linguistic information is acquired and integrated with existing knowledge is provided by the two-stage ‘complementary systems account’ of language learning (McClelland, 2013). Under this view, short-term memory for recently encountered stimuli is primarily supported by hippocampal memory systems, and that sleep plays a key role in the subsequent offline consolidation of this knowledge by allowing these hippocampally mediated memories to be replayed and strengthened within cortical networks (Davis & Gaskell, 2009). It has been suggested that this two-stage approach to long-term acquisition of knowledge about words facilitates carefully integration of old and new information and prevents unhelpful overwriting of existing knowledge (Mirkovic & Gaskell, 2016). Critically, this view suggests that not all learning about word meanings necessarily reflects direct changes to the lexicon itself.

Current evidence in support of the complementary learning systems (CLS) model of word learning comes almost exclusively from experiments in which entirely new wordforms are learned (e.g., Dumay & Gaskell, 2007; Gaskell & Dumay, 2003). However recent evidence suggests that similar mechanisms may support learning about word meanings. Gaskell et al., (2018) found that the magnitude of word-meaning priming seen after 12 or 24 hours is significantly boosted when the learning takes place immediately before a period of sleep. This suggests that the short-term boost in word meaning availability that is seen at intervals of 20-40 minutes (e.g., Rodd et al., 2016) may not reflect a long-term resilient change to the stored mental lexicon, and that it is only with sleep that the knowledge acquired about word-meaning availability is integrated into the lexicon. Specifically, Gaskell et al. (2018) suggested that the hippocampus facilitates the encoding or binding of new associations between the words in the prime sentences, and that this newly acquired hippocampal representation could “form a second source of information alongside more permanent lexical semantic knowledge” that can influence performance on the latter word association test. They suggest that this hippocampal trace would tend to decay within the course of a day, but that sleep can facilitate the use of this newly acquired information to modulate long-term stored (cortical) lexical-semantic knowledge. This account of same-day word meaning priming is inconsistent with earlier claims that word-meaning priming reflects direct modulation of stored (cortical) lexical-semantic knowledge (Gilbert et al., 2018; Rodd et al., 2016, 2013).

Taken together, these studies of word-meaning learning provide a view of lexical learning in which listeners/readers build up new, temporary representations of the ongoing discourse that they encounter. These representations not only provide the immediate disambiguating context that facilitates processing ambiguous words but the binding of words within this discourse to other elements of meaning can alter the disambiguation of ambiguous words that are encountered on the time-scales of 20-40 minutes (Rodd et al., 2016). Importantly, sleep-related consolidation can allow some aspects of these experiences to be retained overall longer time periods. This form of learning ensures that the complex high-dimensional lexical-semantic space is continually being adjusted, maintained and refined by making small, incremental changes such that the future cognitive load associated with understanding ambiguous word in context is minimised (Johnsrude & Rodd, 2016).

Preliminary evidence has suggested that the extent to which lexical-semantic space is reshaped reduces over the lifespan with older adults weighting their recent experiences less heavily than their younger peers (Rodd et al., 2016), but these data come from naturalistic correlational studies and more carefully constructed lab-based experiments are needed to confirm this finding and explore potential explanations.

Future Directions: Towards a Neurobiological Account of Word-Meaning Access

This article has focused providing a purely *cognitive* description of how word meanings are learned, represented and processed. Future work is clearly needed to link this theoretical framework with *neurobiological* accounts that specify more precisely the neuroanatomical basis of (i) stored lexical-semantic knowledge and (ii) any additional cognitive resources that support fluent disambiguation. Numerous fMRI studies of lexical ambiguity have attempted to provide the experimental data that is needed to constrain such neurobiological models. These studies have typically compared high-ambiguity sentences with low-ambiguity control sentences. These studies have consistently identified (i) posterior regions of the middle/inferior temporal lobe and (ii) posterior regions of the left inferior frontal cortex (LIFG) (Bekinschtein, Davis, Rodd, & Owen, 2011; Davis et al., 2007; Rodd, Davis, & Johnsrude, 2005; Rodd, Vitello, Woollams, & Adank, 2015; Rodd, Johnsrude, & Davis, 2012; Rodd, Longe, et al., 2010; Tahmasebi et al., 2012; Vitello, Warren, Devlin, & Rodd, 2014; Zempleni, Renken, Hoeks, Hoogduin, & Stowe, 2007). However the interpretation of these results remains disputed and there remains relative little consensus about the precise functional roles of these brain regions.

There is widespread agreement that brain regions in the left temporal cortex (and inferior parietal cortex) provide a range of stored lexical knowledge, including the phonological, syntactic and semantic properties of words (Hagoort, 2005, 2014). It might therefore seem plausible that the ambiguity-related activations that have been seen in the posterior *temporal* lobe might reflect increased processing that is required within stored lexical-semantic knowledge to settle on the meanings of ambiguous words. However, it is far from straightforward to assume that this region is a key ‘storage’ area for lexical-semantic knowledge. Some authors attribute this region a role in providing the ‘sound to meaning mapping’ that is required for spoken words (Hickok & Poeppel, 2004, 2007, 2015). Under this view this region is not the site for storage of amodal lexical-semantic knowledge, instead, it is hypothesized to store information regarding “the relation (or correspondences) between phonologic information on the one hand and conceptual information on the other” (Hickok & Poeppel, 2015). Other researchers have argued that the posterior regions of the middle temporal lobe supports the control processes that operate on semantic knowledge that is assumed to be stored in more anterior regions of the temporal lobe (Jefferies, 2013; Jefferies & Lambon Ralph, 2006; Noonan, Jefferies, Visser, & Lambon Ralph, 2013; Ralph, Jefferies, Patterson, & Rogers, 2016; Thompson et al., 2018). It is therefore premature to make strong claims about the neurobiological substrate of the distributed lexical-semantic knowledge that is accessed in response to spoken and written words, except to say that these representations are likely to be (i) widely distributed across brain regions, (ii) incorporate diverse aspects of meaning, and (iii) should *not* be viewed as a discrete, localised ‘module’ of the brain.

A similar lack of consensus is seen with respect to the frontal lobe activations that are seen in response to high-ambiguity sentences. Some of the first evidence for the involvement in frontal regions in lexical-semantic disambiguation came from studies showing that patients who had been classified as Broca’s Aphasics had difficulties in using contextual cues when retrieving subordinate word meanings (Hagoort, 1993; Swaab, Brown, & Hagoort, 1998; Swinney, Zurif, & Nicol, 1989). Convergent, and more anatomically precise, evidence also emerged from neuroimaging studies of sentence processing: fMRI studies have very consistently shown activation in the posterior portion of the left inferior frontal gyrus (LIFG; sometimes referred to as Broca’s Area) in response to sentences containing ambiguous words, compared with low-ambiguity sentences (Bekinschtein, Davis, Rodd, & Owen, 2011; Davis et al., 2007; Rodd, Davis, & Johnsrude, 2005; Rodd, Vitello, Woollams, & Adank, 2015; Rodd, Johnsrude, & Davis, 2012; Rodd, Longe, et al., 2010; Tahmasebi et al., 2012; Vitello, Warren, Devlin, & Rodd, 2014; Zempleni, Renken, Hoeks, Hoogduin, & Stowe, 2007).

There is general agreement in the field that although these frontal regions are critical for fluent processing of lexical-semantic knowledge, they do not themselves contain stored lexical-semantic knowledge. Neuropsychological evidence is strongly in support of this view. In particular, patients with semantic aphasia (SA), who typically have large left-hemisphere lesions showing maximal overlap in left inferior frontal gyrus (LIFG) appear to have difficulty accessing semantic knowledge in a flexible and task-appropriate way, despite their store of semantic information being largely intact (Corbett, Jefferies, & Lambon Ralph, 2011; Jefferies & Lambon Ralph, 2006; Noonan et al., 2013; Rogers, Patterson, Jefferies, & Lambon Ralph, 2015). This view

that the LIFG is important for guiding or controlling access to lexical semantic knowledge is consistent with a wealth of evidence from fMRI showing elevated responses when the tasks or stimuli place increased demands on cognitive control systems that have described the contribution of the LIFG in terms of 'semantic selection' or 'controlled retrieval' (e.g., Thompson-Schill, D'Esposito, Aguirre, & Farah, 1997; Wagner, Paré-Blagoev, Clark, & Poldrack, 2001). This view that Broca's area's contribution to lexical-semantic disambiguation should be viewed in terms of providing cognitive control and not stored semantic knowledge is also consistent with evidence that the role of this region in sentence processing is NOT restricted to the processing of *semantic* knowledge, but also shows elevated responses when sentences' complexity/ambiguity arises due to their syntactic structure (e.g., Rodd, Longe, Randall, & Tyler, 2010; Rodd, Vitello, Woollams, & Adank, 2015). Finally, evidence from dual-task methodologies also indicate the involvement of relatively general cognitive mechanisms to lexical-semantic disambiguation (Rodd, Johnsrude, & Davis, 2010).

This view of LIFG function is consistent with the view that with respect to semantic disambiguation, the contribution of the LIFG should be viewed as a top-down control/selection mechanism that modulates stored lexical-semantic knowledge. Under this view the extent to which different stored word meanings become active can depend not only on bottom-up information being received via current perceptual signals, but is also influenced by direct top-down modulation of this knowledge guided by higher level sentence, discourse or situational information. Under this view, frontal control systems provide vital error signals and top-down control mechanisms that guide lexical semantic access and prevent the system becoming 'stuck' in inappropriate settled meaning states. The requirement for top-down modulation is assumed to be increased for high-ambiguity sentences due to the increased likelihood that the default/preferred interpretation of the ambiguous word will be incorrect.

However this is not the only theoretical framework that can provide an explanation for the involvement of the LIFG in semantic disambiguation. An alternative view is the LIFG provides the workspace within which complex sentence and discourse level representations are built bottom-up from incoming lexical semantic knowledge. Under this view, lexical semantic information is delivered to the LIFG from activated lexical-semantic knowledge and that the LIFG provides the necessary workspace to construct complex, combinatorial higher-order sentence/discourse level representations. For example, Hagoort has suggested that this region provides a 'unification space' in which elements of lexical-semantic meaning are unified in to higher-order representations (Hagoort, 2005; Hagoort, 2014). Under this view, the increased activation in the LIFG in response to highly ambiguous sentences reflects the increased unification demands that arise when conflicting elements of lexical knowledge are accessed in response to ambiguous words. It is the unification process that supports true disambiguation as distinct elements of the sentence are integrated into a coherent whole.

These two accounts of the role of the LIFG have proved to be difficult to distinguish using experimental paradigms and further research is needed to characterise more precisely the role of the region in dealing with lexical-semantic ambiguities, and how the functional capacity that is provided by this region interacts with stored (temporal lobe- lexical –semantic knowledge) to support rapid fluent comprehension.

Conclusions

In this review I have integrated findings from a range of experimental paradigms to provide an account of how readers and listeners are able to fluently and rapidly understand word meanings. I have suggested that complex, high-dimensional lexical-semantic representations provide the foundation for semantic disambiguation. The attractor structure of this space ensures that familiar word meanings correspond to stable settled states and prevent the network from settling into unhelpful blend states that contain incompatible elements of meaning. Learning mechanisms play a key role in shaping and maintaining these high quality lexical-semantic representations throughout the lifespan that allow individuals to (i) adapt to changing linguistic environments and (ii) make use of the very recent experience to increase the availability of those meanings that are most likely to be encountered in the near future. Importantly, these learning mechanisms provide the necessary sensitivity to the range of linguistic and non-linguistic cues that can support fluent disambiguation. Finally, I have highlighted the need for future work to link this purely *cognitive* account with current neurobiological accounts of language processing.

References

- Armstrong, B. C., & Plaut, D. C. (2008). Settling dynamics in distributed networks explain task differences in semantic ambiguity effects: Computational and behavioural evidence. In B. C. Love, K. McRae, & V. M. Sloutsky (Eds.), *Proceedings of the 30th Annual Meeting of the Cognitive Science Society* (pp. 273–278). Austin, TX.
- Armstrong, B. C., & Plaut, D. C. (2016). Disparate semantic ambiguity effects from semantic processing dynamics rather than qualitative task differences. *Language, Cognition and Neuroscience*, 31(7), 940–966. <https://doi.org/10.1080/23273798.2016.1171366>
- Bekinschtein, T. A., Davis, M. H., Rodd, J. M., & Owen, A. M. (2011). Why clowns taste funny: The relationship between humor and semantic ambiguity. *Journal of Neuroscience*, 31(26). <https://doi.org/10.1523/JNEUROSCI.5058-10.2011>
- Betts, H. N., Gilbert, R. A., Cai, Z. G., Okedara, Z. B., & Rodd, J. M. (2017). Retuning of Lexical-Semantic Representations: Repetition and Spacing Effects in Word-Meaning Priming. *Journal of Experimental Psychology: Learning Memory and Cognition*, 44(7), 1130–1150. <https://doi.org/10.1037/xlm0000507>
- Borowsky, R., & Masson, M. E. J. (1996). Semantic ambiguity effects in word identification. *Journal of Experimental Psychology: Learning Memory and Cognition*, 22(1).
- Cai, Z. G., Gilbert, R. A., Davis, M. H., Gaskell, M. G., Farrar, L., Adler, S., & Rodd, J. M. (2017). Accent modulates access to word meaning: Evidence for a speaker-model account of spoken word recognition. *Cognitive Psychology*, 98, 73–101. <https://doi.org/10.1016/j.cogpsych.2017.08.003>
- Clark, H. H., & Gerrig, R. J. (1983). Understanding old words with new meanings. *Journal of Verbal Learning and Verbal Behavior*, 22(5). [https://doi.org/10.1016/S0022-5371\(83\)90364-X](https://doi.org/10.1016/S0022-5371(83)90364-X)
- Copestake, A., & Briscoe, T. (1995). Semi-productive Polysemy and Sense Extension. *Journal of Semantics*, 12(1), 15–67. <https://doi.org/10.1093/jos/12.1.15>
- Corbett, F., Jefferies, E., & Lambon Ralph, M. A. (2011). Deregulated semantic cognition follows prefrontal and temporo-parietal damage: Evidence from the impact of task constraint on nonverbal object use. *Journal of Cognitive Neuroscience*, 23(5), 1125–1135. <https://doi.org/10.1162/jocn.2010.21539>
- Davis, M. H., Coleman, M. R., Absalom, A. R., Rodd, J. M., Johnsrude, I. S., Matta, B. F., ... Menon, D. K. (2007). Dissociating speech perception and comprehension at reduced levels of awareness. *Proceedings of the National Academy of Sciences of the United States of America*, 104(41). <https://doi.org/10.1073/pnas.0701309104>
- Davis, M. H., & Gaskell, M. G. (2009). A complementary systems account of word learning: Neural and behavioural evidence. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364(1536). <https://doi.org/10.1098/rstb.2009.0111>
- Duffy, S. A., Kambe, G., & Rayner, K. (2001). The effect of prior disambiguating context on the comprehension of ambiguous words: Evidence from eye movements. In *Decade of behavior. On the consequences of meaning selection: Perspectives on resolving lexical ambiguity* (pp. 27–43). Washington: American Psychological Association. <https://doi.org/10.1037/10459-002>
- Duffy, S. A., Morris, R. K., & Rayner, K. (1988). Lexical ambiguity and fixation times in reading. *Journal of Memory and Language*, 27(4), 429–446. [https://doi.org/10.1016/0749-596X\(88\)90066-6](https://doi.org/10.1016/0749-596X(88)90066-6)
- Dumay, N., & Gaskell, M. G. (2007). Sleep-associated changes in the mental representation of spoken words: Research report. *Psychological Science*, 18(1). <https://doi.org/10.1111/j.1467-9280.2007.01845.x>
- Fang, X., Perfetti, C., & Stafura, J. (2017). Learning new meanings for known words: biphasic effects of prior knowledge. *Language, Cognition and Neuroscience*, 32(5). <https://doi.org/10.1080/23273798.2016.1252050>
- Fodor, J. A. (1983). *The Modularity of Mind*. Cambridge, MA: MIT Press.
- Frisson, S., & Pickering, M. (2007). The processing of familiar and novel senses of a word: Why reading Dickens is easy but reading Needham can be hard. *Language and Cognitive Processes*, 22(4). <https://doi.org/10.1080/01690960601017013>
- Gaskell, M. G., Cairney, S. A., & Rodd, J. M. (2019). Contextual priming of word meanings is stabilized over sleep. *Cognition*, 182, 109–126. <https://doi.org/10.1016/j.cognition.2018.09.007>
- Gaskell, M. G., & Dumay, N. (2003). Lexical competition and the acquisition of novel words. *Cognition*, 89(2). [https://doi.org/10.1016/S0010-0277\(03\)00070-2](https://doi.org/10.1016/S0010-0277(03)00070-2)

- Gilbert, R. A., Davis, M. H., Gaskell, M. G., & Rodd, J. M. (2018). Listeners and Readers Generalize Their Experience With Word Meanings Across Modalities. *Journal of Experimental Psychology: Learning Memory and Cognition*. <https://doi.org/10.1037/xlm0000532>
- Hagoort, P. (1993). Impairments of lexical-semantic processing in aphasia: Evidence from the processing of lexical ambiguities. *Brain and Language*, 45(2), 189–232. <https://doi.org/10.1006/brln.1993.1043>
- Hagoort, P. (2005). On Broca, brain, and binding: A new framework. *Trends in Cognitive Sciences*, 9(9). <https://doi.org/10.1016/j.tics.2005.07.004>
- Hagoort, P. Nodes and networks in the neural architecture for language: Broca's region and beyond, 28 *Current Opinion in Neurobiology* § (2014). <https://doi.org/10.1016/j.conb.2014.07.013>
- Hickok, G., & Poeppel, D. (2004). Dorsal and ventral streams: A framework for understanding aspects of the functional anatomy of language. *Cognition*, 92(1–2). <https://doi.org/10.1016/j.cognition.2003.10.011>
- Hickok, G., & Poeppel, D. (2007). The cortical organization of speech processing. *Nature Reviews Neuroscience*, 8(5). <https://doi.org/10.1038/nrn2113>
- Hickok, G., & Poeppel, D. Neural basis of speech perception, 129 § (2015). <https://doi.org/10.1016/B978-0-444-62630-1.00008-1>
- Hogaboam, T. W., & Perfetti, C. A. (1975). Lexical ambiguity and sentence comprehension. *Journal of Verbal Learning and Verbal Behavior*, 14(3), 265–274. [https://doi.org/10.1016/S0022-5371\(75\)80070-3](https://doi.org/10.1016/S0022-5371(75)80070-3)
- Hulme, R. C., Barsky, D., & Rodd, J. M. (2018). Incidental learning and long-term retention of new word meanings from stories: The effect of number of exposures. *In Press at Language Learning*.
- Jefferies, E. (2013). The neural basis of semantic cognition: Converging evidence from neuropsychology, neuroimaging and TMS. *Cortex*, 49(3), 611–625. <https://doi.org/10.1016/j.cortex.2012.10.008>
- Jefferies, E., & Lambon Ralph, M. A. (2006). Semantic impairment in stroke aphasia versus semantic dementia: A case-series comparison. *Brain*, 129(8), 2132–2147. <https://doi.org/10.1093/brain/awl1153>
- Johnsrude, I. S., & Rodd, J. M. (2016). Factors That Increase Processing Demands When Listening to Speech. *Neurobiology of Language*, 491–502. <https://doi.org/10.1016/B978-0-12-407794-2.00040-7>
- Joordens, S., & Besner, D. (1994). When Banking on Meaning Is Not (Yet) Money in the Bank: Explorations in Connectionist Modeling. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20(5).
- Kambe, G., Rayner, K., & Duffy, S. A. (2001). Global context effects on processing lexically ambiguous words: Evidence from eye fixations. *Memory and Cognition*, 29(2).
- Kawamoto, A. H., Farrar IV, W. T., & Kello, C. T. (1994). When Two Meanings Are Better Than One: Modeling the Ambiguity Advantage Using a Recurrent Distributed Network. *Journal of Experimental Psychology: Human Perception and Performance*, 20(6).
- Kellas, G., & Vu, H. (1999). Strength of context does modulate the subordinate bias effect: a reply to Binder and Rayner. *Psychonomic Bulletin & Review*, 6(3), 511–517; discussion 518.
- Klein, D. E., & Murphy, G. L. (2001). The Representation of Polysemous Words. *Journal of Memory and Language*, 45(2). <https://doi.org/10.1006/jmla.2001.2779>
- McClelland, J. L. (2013). Incorporating rapid neocortical learning of new schema-consistent information into complementary learning systems theory. *Journal of Experimental Psychology: General*, 142(4), 1190–1210. <https://doi.org/10.1037/a0033812>
- Mirkovic, J., & Gaskell, M. G. (2016). Does sleep improve your grammar? Preferential consolidation of arbitrary components of new linguistic knowledge. *PLoS ONE*, 11(4). <https://doi.org/10.1371/journal.pone.0152489>
- Noonan, K. A., Jefferies, E., Visser, M., & Lambon Ralph, M. A. (2013). Going beyond inferior prefrontal involvement in semantic control: evidence for the additional contribution of dorsal angular gyrus and posterior middle temporal cortex. *Journal of Cognitive Neuroscience*, 25(11), 1824–1850. https://doi.org/10.1162/jocn_a_00442
- Onifer, W., & Swinney, D. A. (1981). Accessing lexical ambiguities during sentence comprehension: Effects of frequency of meaning and contextual bias. *Memory & Cognition*, 9(3). <https://doi.org/10.3758/BF03196957>
- Pacht, J. M., & Rayner, K. (1993). The processing of homophonic homographs during reading: Evidence from eye movement studies. *Journal of Psycholinguistic Research*, 22(2), 251–271. <https://doi.org/10.1007/BF01067833>
- Piantadosi, S. T., Tily, H., & Gibson, E. (2012). The communicative function of ambiguity in language. *Cognition*, 122(3), 280–291. <https://doi.org/10.1016/J.COGNITION.2011.10.004>

- Plaut, D. C., McClelland, J. L., Seidenberg, M. S., & Patterson, K. (1996). Understanding Normal and Impaired Word Reading: Computational Principles in Quasi-Regular Domains. *Psychological Review*, 103(1).
- Rabagliati, H., Dumas, L. A. A., & Bemis, D. K. (2017). Representing composed meanings through temporal binding. *Cognition*, 162, 61–72. <https://doi.org/10.1016/j.cognition.2017.01.013>
- Ralph, M. A. L., Jefferies, E., Patterson, K., & Rogers, T. T. (2016). The neural and computational bases of semantic cognition. *Nature Reviews Neuroscience*, 18(1), 42–55. <https://doi.org/10.1038/nrn.2016.150>
- Rayner, K., Binder, K. S., & Duffy, S. A. (1999). Contextual Strength and the Subordinate Bias Effect: Comment on Martin, Vu, Kellas, and Metcalf. *Quarterly Journal of Experimental Psychology Section A: Human Experimental Psychology*, 52(4), 841–852. <https://doi.org/10.1080/713755868>
- Rayner, K., Pacht, J. M., & Duffy, S. A. (1994). Effects of Prior Encounter and Global Discourse Bias on the Processing of Lexically Ambiguous Words: Evidence From Eye Fixations. *Journal of Memory and Language*, 33(4), 527–544. <https://doi.org/10.1006/jmla.1994.1025>
- Rodd, J. M., Berriman, R., Landau, M., Lee, T., Ho, C., Gaskell, M. G., & Davis, M. H. (2012). Learning new meanings for old words: Effects of semantic relatedness. *Memory and Cognition*, 40(7). <https://doi.org/10.3758/s13421-012-0209-1>
- Rodd, J. M., Cai, Z. G., Betts, H. N., Hanby, B., Hutchinson, C., & Adler, A. (2016). The impact of recent and long-term experience on access to word meanings: Evidence from large-scale internet-based experiments. *Journal of Memory and Language*, 87. <https://doi.org/10.1016/j.jml.2015.10.006>
- Rodd, J. M., Davis, M. H., & Johnsrude, I. S. (2005). The neural mechanisms of speech comprehension: fMRI studies of semantic ambiguity. *Cerebral Cortex*, 15(8). <https://doi.org/10.1093/cercor/bhi009>
- Rodd, J. M., Gaskell, G., & Marslen-Wilson, W. (2002). Making sense of semantic ambiguity: Semantic competition in lexical access. *Journal of Memory and Language*, 46(2). <https://doi.org/10.1006/jmla.2001.2810>
- Rodd, J. M., Gaskell, M. G., & Marslen-Wilson, W. D. (2004). Modelling the effects of semantic ambiguity in word recognition. *Cognitive Science*, 28(1). <https://doi.org/10.1016/j.cogsci.2003.08.002>
- Rodd, J. M., Johnsrude, I. S., & Davis, M. H. (2010). The role of domain-general frontal systems in language comprehension: Evidence from dual-task interference and semantic ambiguity. *Brain and Language*, 115(3). <https://doi.org/10.1016/j.bandl.2010.07.005>
- Rodd, J. M., Johnsrude, I. S., & Davis, M. H. (2012). Dissociating frontotemporal contributions to semantic ambiguity resolution in spoken sentences. *Cerebral Cortex*, 22(8), 1761–1773. <https://doi.org/10.1093/cercor/bhr252>
- Rodd, J. M., Longe, O. A., Randall, B., & Tyler, L. K. (2010). The functional organisation of the fronto-temporal language system: Evidence from syntactic and semantic ambiguity. *Neuropsychologia*, 48(5). <https://doi.org/10.1016/j.neuropsychologia.2009.12.035>
- Rodd, J. M., Lopez Cutrin, B., Kirsch, H., Millar, A., & Davis, M. H. (2013). Long-term priming of the meanings of ambiguous words. *Journal of Memory and Language*, 68(2). <https://doi.org/10.1016/j.jml.2012.08.002>
- Rodd, J. M., Vitello, S., Woollams, A. M., & Adank, P. (2015). Localising semantic and syntactic processing in spoken and written language comprehension: An Activation Likelihood Estimation meta-analysis. *Brain and Language*, 141. <https://doi.org/10.1016/j.bandl.2014.11.012>
- Rogers, T. T., Patterson, K., Jefferies, E., & Lambon Ralph, M. A. (2015). Disorders of representation and control in semantic cognition: Effects of familiarity, typicality, and specificity. *Neuropsychologia*, 76, 220–239. <https://doi.org/10.1016/j.neuropsychologia.2015.04.015>
- Seidenberg, M. S., Tanenhaus, M. K., Leiman, J. M., & Bienkowski, M. (1982). Automatic access of the meanings of ambiguous words in context: Some limitations of knowledge-based processing. *Cognitive Psychology*, 14(4), 489–537. [https://doi.org/10.1016/0010-0285\(82\)90017-2](https://doi.org/10.1016/0010-0285(82)90017-2)
- Sereno, S. C. (1995). Resolution of Lexical Ambiguity: Evidence From an Eye Movement Priming Paradigm. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21(3), 582–595. <https://doi.org/10.1037/0278-7393.21.3.582>
- Sheridan, H., Reingold, E. M., & Daneman, M. (2009). Using puns to study contextual influences on lexical ambiguity resolution: Evidence from eye movements. *Psychonomic Bulletin and Review*, 16(5), 875–881. <https://doi.org/10.3758/PBR.16.5.875>
- Srinivasan, M., & Rabagliati, H. (2015). How concepts and conventions structure the lexicon: Cross-linguistic evidence from polysemy. *Lingua*, 157, 124–152. <https://doi.org/10.1016/j.lingua.2014.12.004>

- Swaab, T. Y., Brown, C., & Hagoort, P. (1998). Understanding ambiguous words in sentence contexts: Electrophysiological evidence for delayed contextual selection in Broca's aphasia. *Neuropsychologia*, 36(8), 737–761. [https://doi.org/10.1016/S0028-3932\(97\)00174-7](https://doi.org/10.1016/S0028-3932(97)00174-7)
- Swinney, D. A. (1979). Lexical access during sentence comprehension: (Re)consideration of context effects. *Journal of Verbal Learning and Verbal Behavior*, 18(6), 645–659. [https://doi.org/10.1016/S0022-5371\(79\)90355-4](https://doi.org/10.1016/S0022-5371(79)90355-4)
- Swinney, D. A., Zurif, E., & Nicol, J. (1989). The effects of focal brain damage on sentence processing: An examination of the neurological organization of a mental module. *Journal of Cognitive Neuroscience*, 1(1), 25–37. <https://doi.org/10.1162/jocn.1989.1.1.25>
- Tabossi, P. (1988). Accessing lexical ambiguity in different types of sentential contexts. *Journal of Memory and Language*, 27(3), 324–340. [https://doi.org/10.1016/0749-596X\(88\)90058-7](https://doi.org/10.1016/0749-596X(88)90058-7)
- Tabossi, P., Colombo, L., & Job, R. (1987). Accessing lexical ambiguity: Effects of context and dominance. *Psychological Research*, 49(2–3), 161–167. <https://doi.org/10.1007/BF00308682>
- Tahmasebi, A. M., Davis, M. H., Wild, C. J., Rodd, J. M., Hakyemez, H., Abolmaesumi, P., & Johnsrude, I. S. (2012). Is the link between anatomical structure and function equally strong at all cognitive levels of processing? *Cerebral Cortex*, 22(7). <https://doi.org/10.1093/cercor/bhr205>
- Thompson-Schill, S. L., D'Esposito, M., Aguirre, G. K., & Farah, M. J. (1997). Role of left inferior prefrontal cortex in retrieval of semantic knowledge: A reevaluation. *Proceedings of the National Academy of Sciences of the United States of America*, 94(26). <https://doi.org/10.1073/pnas.94.26.14792>
- Thompson, H. E., Almaghyuli, A., Noonan, K. A., barak, O., Lambon Ralph, M. A., & Jefferies, E. (2018). The contribution of executive control to semantic cognition: Convergent evidence from semantic aphasia and executive dysfunction. *Journal of Neuropsychology*, 12(2), 312–340. <https://doi.org/10.1111/jnp.12142>
- Twilley, L. C., Dixon, P., Taylor, D., & Clark, K. (1994). University of Alberta norms of relative meaning frequency for 566 homographs. *Memory & Cognition*, 22(1), 111–126. <https://doi.org/10.3758/BF03202766>
- Vitello, S. (2013). *Cognitive and neural mechanisms underlying semantic ambiguity resolution*. PhD Thesis, University College London.
- Vitello, S., & Rodd, J. M. (2015). Resolving Semantic Ambiguities in Sentences: Cognitive Processes and Brain Mechanisms. *Linguistics and Language Compass*, 9(10). <https://doi.org/10.1111/lnc3.12160>
- Vitello, S., Warren, J. E., Devlin, J. T., & Rodd, J. M. (2014). Roles of frontal and temporal regions in reinterpreting semantically ambiguous sentences. *Frontiers in Human Neuroscience*, 8(JULY). <https://doi.org/10.3389/fnhum.2014.00530>
- Wagner, A. D., Paré-Blagoev, E. J., Clark, J., & Poldrack, R. A. (2001). Recovering meaning: Left prefrontal cortex guides controlled semantic retrieval. *Neuron*, 31(2). [https://doi.org/10.1016/S0896-6273\(01\)00359-2](https://doi.org/10.1016/S0896-6273(01)00359-2)
- Wiley, J., George, T., & Rayner, K. (2016). Baseball fans don't like lumpy batters: Influence of domain knowledge on the access of subordinate meanings. *Quarterly Journal of Experimental Psychology*, 1–11. <https://doi.org/10.1080/17470218.2016.1251470>
- Witzel, J., & Forster, K. (2015). Lexical co-occurrence and ambiguity resolution. *Language, Cognition and Neuroscience*, 29(2), 158–185. <https://doi.org/10.1080/01690965.2012.748925>
- Zemleni, M.-Z., Renken, R., Hoeks, J. C. J., Hoogduin, J. M., & Stowe, L. A. (2007). Semantic ambiguity processing in sentence context: Evidence from event-related fMRI. *NeuroImage*, 34(3). <https://doi.org/10.1016/j.neuroimage.2006.09.048>