Word meaning access: The one-to-many mapping from form to meaning

Jennifer M. Rodd

Department of Experimental Psychology, University College London

To Appear in:

The OUP Handbook of the Mental Lexicon

THIS IS A DRAFT VERSION OF A BOOK CHAPTER.

IT HAS NOT YET UNDERGONE PEER REVIEW.

Abstract

This chapter focuses on the process by which stored knowledge about a word's form (orthographic or phonological) maps onto stored knowledge about its meaning. This mapping is made challenging by the ambiguity that is ubiquitous in natural language: most familiar words can refer to multiple different concepts. This one-to-many mapping from form to meaning within the lexicon is a core feature of word-meaning access. Fluent, accurate word-meaning access requires that comprehenders integrate multiple cues in order to determine which of a word's possible semantic features are relevant in the current context. Specifically, word-meaning access is guided by (i) distributional information about the a priori relative likelihoods of different word meanings and (ii) a wide range of contextual cues that indicate which meanings are most likely in the current context.

Words: The Gateway to Meaning

The mental lexicon contains stored knowledge about familiar words: information about their spoken and written forms, their meanings and their grammatical properties (see other chapters in this volume). A complete theoretical account of lexical processing must specify not only the nature of these different forms of stored, lexical knowledge, but must also describe the computational processes by which these different forms of knowledge interact. Here we focus on the process by which stored knowledge about a word's form (either orthographic or phonological) maps forward onto knowledge about its meaning.

Wordforms are the gateway to meaning: once a reader (or listener) has identified the particular spoken or printed word that is present in their environment they are able to access a rich store of conceptual knowledge that may include perceptual, abstract and episodic features of the relevant concept. These individual word-sized units of stored knowledge can be combined together in a hierarchical manner within sentences to communicate complex, novel ideas.

A common approach in models of lexical processing is to assume that the lexicon contains abstract representations of individual words. For example, in Morton's highly influential *Logogen* model (1969), any given logogen becomes activated by information (perceptual or contextual) that is consistent with the properties of that specific word. Once a single logogen's activation level rises to its threshold level, the properties of that word (including its meaning) becomes available to the comprehender. This idea that individual words are represented by abstract lexical units is retained in subsequent models such as the *Cohort* model of spoken word recognition (Marslen-Wilson & Welsh, 1978), the *Interactive Activation and Competition* model of visual word recognition (McClelland & Rumelhart, 1981), and the *Dual Route Cascaded* model of reading aloud (Coltheart, Rastle, Perry, Langdon, & Ziegler, 2001).

Thus, according to these models, lexical representations provide an abstract level of lexical knowledge that bridges the gap between form-based (phonological or orthographic) representations and knowledge about word meanings.

Distributed connectionist models of reading provide an alternative framework for viewing this mapping from form to meaning (Plaut, McClelland, Seidenberg, & Patterson, 1996; Seidenberg & McClelland, 1989; see Figure 1). This ‘triangle’ model of lexical processing proposes that three distinct sets of lexical knowledge (orthographic, phonological, and semantic) words are each represented as distributed patterns of activity over three separate featural representations. Importantly, this theoretical approach does not assume that there are individual, localist units that each correspond to particular, familiar words, instead knowledge about individual words is coded in a distributed manner as patterns of activation over an array of units. Thus, within this framework, there is a direct mapping from information about a word’s form to its meaning without the need for intervening abstract level of lexical knowledge. Different lexical tasks (such as reading aloud) are viewed as transformations among these representations. For example, Seidenberg & McClelland (1989) reported simulations that focus on reading aloud: the model is presented with an input orthographic pattern and is required to learn how to generate an appropriate output phonological pattern. Although these early connectionist models did not fully implement the word-meaning representations, and focused primarily on the mapping from orthography to phonology, this framework has provided the foundation for much of the computational modelling work on word-meaning access in the past 30 years (see Rodd, 2019 for a recent review).

Figure 1. Seidenberg and McClelland's (1989) lexical processing framework. Each oval represents a group of units, and each arrow represents a group of connections between these units. The unlabeled ovals represent 'hidden' units that are included in the model to aid its learning of the complex mappings between different forms of lexical knowledge. Adapted from Figure 1, "A Distributed, Developmental Model of Word Recognition and Naming," by M. S. Seidenberg and J. L. McClelland (1989), *Psychological Review*, 96(4).

This computational modelling work highlighted the arbitrary nature of the mapping from form-to-meaning. For mono-morphemic words such as “dog” and “cat, the individual constituent letters/sounds do not provide strong statistical cues as to the meanings of these words: words

with similar sounds do not always have similar meanings. This arbitrariness in the form-to-meaning mapping is in stark contrast to the quasi-regular mapping between words' printed forms and spoken forms in English (and many other alphabetic languages): when presented with a novel wordform such as 'wug' readers will have highly consistent intuitions about how this word should be pronounced, but do not have similar certainty about its likely meaning (Castles, Rastle, & Nation, 2018). Thus all theoretical accounts of how the form-to-meaning mapping is computed in human comprehenders must provide a mechanisms that is capable of learning such an arbitrary mapping. But subsequent computational modelling work that focused in more detail on the form-to-meaning mapping (Armstrong & Plaut, 2008; Borowsky & Masson, 1996; Joordens & Besner, 1994; Kawamoto, Farrar, & Kello, 1994; Rodd, Gaskell, & Marslen-Wilson, 2004) highlighted a second, more problematic characteristic of this mapping: knowing uniquely exactly which wordform has been seen or heard (i.e. the word was definitely "bark" and not "park") does not uniquely identify a single word meaning. For example the word "bark" might be used to refer to the outer covering of a tree, but might also be used to refer to the sound made by a dog.

This chapter will argue that the mapping from form to meaning is best characterized as a *one-to-many* mapping: most words in language are, to some extent, ambiguous, such that a single wordform maps forward onto multiple, alternative word meanings. This high level of ambiguity places important challenges on theories of lexical processing: knowing which form of a word has been encountered does not, in isolation, provide sufficient information to uniquely identify the word meaning that was intended by the speaker or writer. Comprehenders must therefore use additional constraining information to achieve successful word meaning access. This computational challenge arises both for theories of lexical processing in which word meanings

are being accessed directly from information about word forms (e.g., Seidenberg & McClelland, 1989) and those that assume an intervening level of abstract lexical representation (e.g., Morton, 1969).

The aims of this chapter are therefore to (i) summarize the different forms of lexical-semantic ambiguity that lead to a one-to-many mapping from form to meaning and then (ii) describe what is known from experimental evidence about the constraints that support fluent, accurate word meaning access by guiding comprehenders towards appropriate word meanings. These constraints fall into two broad categories. First, listeners and readers make sophisticated use of their recent and longer-term experience with word meanings, such that individual meanings become more readily available as a consequence of experience. Second, a wide range of contextual cues guide access towards those word meanings that are most likely in any given context. Finally, the chapter will review models that specify potential mechanisms by which these different constraints operate to ensure successful word-meaning access.

Forms of lexical-semantic ambiguity

(i) Homonymy

The one-to-many mapping from form to meaning can arise from various forms of lexical-semantic ambiguity. The most salient form of ambiguity occurs for those wordforms that have multiple, different meanings that are not related in either their meaning or their etymological history. For example, it is a historical accident that English happens to use the same wordform “bark” to refer both to the sound made by a dog and to the outer covering of a tree. These

unrelated meanings are usually listed as separate entries in dictionaries and are referred to as homonyms, but they can also be referred to as homographs (because they share their spelling) or homophones (because they share their pronunciation). True homonyms, where there is no relationship at all between the different meanings, are relatively rare – Rodd Gaskell & Marslen-Wilson (2002) estimated that only 7% of common English words are homonyms. English also contains homophones that only share their pronunciation (e.g., “waist/waste”, “some/sum”, “for/four”) as well as homographs that only share their spelling (e.g., “lead”, and “close”). These words, which are also accidents of language evolution (and the imperfect relationship between orthography and phonology in English), are ambiguous only in spoken or written language respectively.

Homonyms present a particular challenge because the ambiguity that is inherent in these words *must* always be fully resolved for the overall sentence to be understood. The semantic features associated with the different meanings are completely non-overlapping and mutually exclusive. Therefore if a comprehender retrieves a mixture of these two meanings (sometimes known as a blend state; Rodd, Gaskell, & Marslen-Wilson, 2004) this will not provide them with a useful, meaningful representation – the comprehender *must* therefore, at some point, fully commit to one or other of the meanings for a single coherent representation to be built. This need to disambiguate homonyms places significant additional processing demands on the language comprehension system that have been observed using a wide range of behavioural and neuroimaging methods (see Johnsrude & Rodd, 2016; Vitello & Rodd, 2015 for reviews).

(ii) Polysemy

A more common form of ambiguity – usually referred to as polysemy – occurs when a single wordform can refer to multiple, semantically related word senses. These word senses often come in relatively large clusters. For example, the word “run” has numerous different dictionary definitions (e.g., “the athlete runs in the race”, “the paint runs down the wall”, “the river runs through the valley”, “the politician runs in the election”, “the programme runs on the computer”, etc.). These different senses often overlap somewhat in their meanings, although it can often be difficult to quantify precisely the nature of these relationships. Indeed, it is likely that the senses of words that are listed in any given dictionary does not even fully capture the full extent of the flexibility of such word, which can often be used in a very wide range of contexts to express a range of subtly different concepts. Therefore, although the presence of polysemy in natural language often increases the processing demands, by making the form-to-meaning mapping more challenging, it adds considerably to the communicative power of language by allowing a finite set of familiar wordforms to convey a far richer array of word meanings than if each wordform mapped uniquely onto a single meaning.

Importantly, the capacity of speakers to productively extend the meaning of familiar wordforms to convey novel, related conceptual ideas provides a key mechanism for language change such that language continues to convey those word meanings that are needed to support communication. For example, “tweet”, “tablet”, “post” and many more words have acquired novel technological meanings that share some meaning overlap with their previous non-technological meanings (Maciejewski, Rodd, Mon-Williams, & Klepousniotou, 2018; Rodd et al., 2012). This creative aspect of polysemy is particularly apparent for ‘regular polysemy’: those clusters of words that have systematic patterns of senses. For example, animal names can be

productively extended to refer to the meat that comes from that animal (e.g., chicken, lamb, ostrich, etc.) and the names of substances can often also refer to objects made from that material (e.g., glass, iron, paper etc; Copestake & Briscoe, 1995; Srinivasan & Rabagliati, 2015). There is evidence from historical analyses of language use that new word senses emerge in a relatively regular manner, predominantly through a process of ‘nearest-neighbor chaining’ whereby each new sense is highly similar in meaning to an existing sense in the language (Ramiro, Srinivasan, Malt, & Xu, 2018).

Polysemy differs from homonymy in several important respects. Most obviously, there is a far greater level of semantic (or associative) overlap between the different meanings. From a processing point of view, the level of competition that results from polysemy can be somewhat reduced compared with homonymy, and the level of competition can vary according to the current task demands (Rodd et al., 2012, 2002). For example, the different senses of ‘run’ are likely to share significant aspects of their meanings so that the only competition that occurs during comprehension will be between those non-shared elements of meaning. In addition, in many task situations the reader may not fully disambiguate between the different senses of a polysemous word– it may be sufficient for them to activate some of the shared/core semantic features of the word. For example, readers can make a lexical decision to the word “run” or decide that it is semantically related to a probe word “move” without necessarily fully disambiguating its word meaning to one precise word sense. In addition, studies of the syntactic aspects of sentence processing have highlighted the extent to which readers often engage in ‘good enough’ processing in which their final interpretations of sentences containing syntactic ambiguities may often be incomplete or inaccurate (Ferreira, Bailey, & Ferraro, 2002). (See

Armstrong & Plaut, 2016; Rodd et al., 2002 for extensive discussion about how the level of semantic disambiguation of polysemous word can vary as a function of task constraints.)

It is important to note that in practice it can be extremely difficult to draw a clear categorical distinction between polysemy and homonymy: it is not always clear exactly how *related* two word meanings need to be to count as polysemous – lexicographers do not always agree on how word meanings/senses should be classified (Rodd et al., 2002). This is important because some theoretical accounts of how word meanings are represented assume a categorical classification in which there is a clear qualitative difference in how these two forms of ambiguity are represented: homonyms correspond to two distinct lexical items with separate representations in the lexicon, while the related sense of polysemous words all correspond to a single lexical entry (Carston, 2019). In contrast, other approaches allow the distinction between related word senses and unrelated word meanings to be viewed as a graded continuum. For example, Rodd (2019) proposed an account of word-meaning access (based on earlier work by Rodd et al., 2004) in which word senses/meanings can vary continuously in the strength of their semantic relationship without need to draw a clear division between these two forms of ambiguity. This latter approach can accommodate findings that suggest that there is considerable variability in relatedness within the class of polysemous words (Klein & Murphy, 2001).

Finally, it is important to note that although lexical-semantic ambiguity is usually viewed in a negative light, as a source of additional potential competition and confusion, it has been argued that, under some conditions, ambiguity is actually a functional property of language that can improve communicative efficiency (Piantadosi, Tily, & Gibson, 2012). Specifically, Piantadosi et al., (2012) have argued that if words are used within rich, informative contexts that provide sufficient cues to allow them to be accurately disambiguated, then re-using a relatively

small set of word forms to convey a larger set of word meanings can reduce the processing demands relative to an alternative system that uses a much larger set of word forms (Piantadosi et al., 2012).

(iii) Contextual Modulation of Unambiguous Word Meanings

Even for relatively unambiguous wordforms that are used by all speakers of a given dialect to refer to a single, clearly defined specific object, action or other conceptual entity, the mapping from form to meaning is not entirely straightforward. For most highly familiar words, individuals will have a large body of stored conceptual knowledge, and the subset of this knowledge that is most relevant can vary according to the specific context in which the word has been used (Yee & Thompson-Schill, 2016). Take for example the word “lemon”. Even if we set aside its lower frequency senses such as referring to a pale yellow colour or to describe something that is unsatisfactory or defective (e.g., “it turned out to be a lemon”) and focus instead on its dominant interpretation as a type of fruit, then the mapping to meaning is somewhat flexible. Despite the likely widespread agreement amongst speakers of English about which objects in the world can (and cannot) be accurately be referred to by the label “lemon”, once this word is placed within a sentence context it is advantageous for the listener to perform some form of context-driven semantic disambiguation. For example, (Tabossi, 1988b) used a cross-modal priming paradigm to demonstrate when relatively unambiguous words like “lemon” are placed within a sentence context, properties like “sourness” are more readily available at the offset of the target word if the preceding context guides the listener towards this particular semantic feature (e.g. “The little boy shuddered after eating a slice of lemon”) compared with sentences where other features are more contextually relevant (e.g. “The little boy was playing

on the floor rolling a lemon”). This demonstrates that context can very rapidly modulate the aspects of a word’s meaning that are accessed/maintained in any given sentence context. (See Mirković & Altmann, 2019 for more recent demonstration of how *visual* context can influence semantic interpretation of unambiguous nouns.)

Taken together, these different forms of lexical-semantic ambiguity require that any theoretical account of word meaning access must specify how readers and listeners are able to select those aspects of lexical meaning most relevant to the current context. Successful comprehension relies on ensuring that the most relevant aspects of meaning are available for integration into a representation of the ongoing discourse. This process of lexical-semantic disambiguation is most salient for homonyms such as “bark”, but is a ubiquitous aspect of lexical processing that is constantly required for successful comprehension.

The following sections will overview what is currently known about how comprehenders deal with this need to select contextually relevant aspects of word meanings. The vast majority of this research has focused on homonym processing: the minimal overlap between the alternative interpretations of these wordforms makes it more straightforward to measure the extent to which meaning(s) have been retrieved by participants in any given experimental scenario. But it is important to keep in mind that lexical-semantic disambiguation is also critical for efficient comprehension of polysemes (e.g., “run”) as well as for relatively unambiguous words. Successful comprehension requires that listeners and readers select from their vast array of stored conceptual knowledge the specific information that is most relevant to the current contextual situation (Yee & Thompson-Schill, 2016).

Cues that Guide Lexical-Semantic Disambiguation

(i) Recent and Longer-Term Experience with Word Meanings

The one-to-many mapping from form to meaning entails that perceptual information alone is rarely sufficient for a reader or listener to be certain about the meaning of each word they encounter. They must use other cues to help them settle on one likely interpretation for each word. One important cue that guides word-meaning access is distributional information about the frequency with which different meanings/senses have previously been encountered: meanings that are more frequent in language (usually referred to as *dominant* meanings) tend to be more readily available (e.g., Hogaboam & Perfetti, 1975). Evidence from a very wide range of behavioural and neuroimaging paradigms have confirmed that more frequent (dominant) word meanings experience a processing benefit that makes them more easily accessible, compared with less frequent (subordinate) meanings (see Vitello & Rodd, 2015 for review).

The clearest evidence that word-meaning dominance strongly modulates word-meaning access comes from the large body of research in which readers' eye-movements are tracked during reading. When readers are presented with text (or any other visual stimulus) they will respond by making a series of discrete eye movements called saccades. Between these saccades, their eyes remain relatively still for short periods of time (usually for 200-300 msec.) as they fixate on a specific area of the text and gain detailed visual information about those letters that are currently in the center of their gaze. Importantly as text becomes conceptually more difficult, readers tend to (i) lengthen duration of their individual fixations, (ii) reduce the distance (i.e. number of letters) that they jump forward on each saccade, and (iii) make more backwards regressions to reread earlier parts of the text (Rayner, 1998). Eye-tracking studies of lexical

ambiguity have consistently shown that when highly dominant word meanings, such as the ‘writing implement’ meaning of “pen”, are preceded by a supportive sentence context then these meanings can be easily accessed with no measurable delay in processing compared to a low-ambiguity control word. In contrast, when the preceding context supports a highly subordinate word meaning, such as the ‘animal enclosure’ meaning of “pen” then there is a delay in reading times that reflects the additional processing demands associated with accessing this less frequent meaning. (See Duffy et al., 2001 for a comprehensive review of findings concerning how ambiguous words are processed within sentence contexts.)

Although eye-tracking studies have been highly influential in driving theoretical accounts of word-meaning access that can account for many aspects of the data seen in eye-tracking studies of lexical ambiguity (e.g., The Re-ordered Access Model: Duffy et al., 2001), one limitation is that they have typically operationalised ‘dominance’ as a single unitary property of word meanings that is assumed to be relatively consistent across participants: dominance is usually estimated by averaging the responses to these words when they are presented in isolation to a *different* group of participants (e.g., Twilley, Dixon, Taylor, & Clark, 1994). Underlying this approach is the assumption that individuals within a particular language community will all have had relatively similar experiences with word meanings, such that a single estimate of meaning dominance can reliably predict behavior across all participants in any given experiment.

More recently however, researchers have attempted to further develop this approach to explore in more detail how the availability of word meanings is influenced by individuals’ idiosyncratic experiences with particular word meanings. Rodd et al. (2016) investigated word meaning preferences in recreational rowers whose language experience differs somewhat from the rest of the population because they regularly encounter rowing-related terminology: they will

frequently encounter rowing-specific meanings for highly familiar words such as “catch”, “square” and “feather”, which refer (for example) to a particular position of the oar or phase of the rowing stroke. Rodd et al., (2016) found that the relative availability of these rowing-specific meanings (compared with the non-rowing, dominant meaning) increased in a logarithmic manner with participants’ number years of rowing experience: the availability of word meanings increase gradually over time in response to an individual’s idiosyncratic linguistic environment over relatively long periods of time. For example, an individual with 10 years’ rowing experience will (on average) have different word-meaning preferences to an individual with 2 years’ experience, even for words for which both of them are highly familiar with the rowing-related meaning. This indicates that individuals are effectively tracking the prevalence of word meanings over relatively long time periods. Similarly, Wiley, George & Rayner (2016) showed that (compared with non-experts) individuals with specific expertise in baseball showed significantly more difficulty processing words that had a baseball-related meaning (e.g., “bat”) when these words were presented in sentence contexts that were 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 data indicate that the availability of the baseball-related meanings had been boosted in the baseball experts as a consequence of their increased experience with these word meanings. These studies emphasise the inevitable differences in how different people represent word meanings – every individual has their own idiosyncratic linguistic history that will shape the distributional statistics with which we encounter the different meanings of words. These studies demonstrate that individual differences in long-term experience directly influence the ease with which different word meanings come to mind. These studies also remind us of the flexible functionality of the language comprehension system. These incremental adjustments of

lexical-semantic knowledge in response to a changing linguistic environment in a manner that will improve future comprehension by boosting the availability of those meanings that they are likely to encounter in the future (Rodd et al., 2016).

In addition to these long-term effects of experience, recent experimental evidence has shown that more encounters with particular, subordinate interpretations of ambiguous words can substantially boost the availability of these primed meanings (Betts, Gilbert, Cai, Okedara, & Rodd, 2018; Gaskell, Cairney, & Rodd, 2019; Gilbert, Davis, Gaskell, & Rodd, 2018; Rodd et al., 2016; Rodd, Lopez Cutrin, Kirsch, Millar, & Davis, 2013). In word-meaning priming experiments, participants encounter a subordinate meaning of an ambiguous word within a strongly disambiguating sentence context (e.g., “The farmer moved the sheep into the PEN”). Then, after a delay (typically 20-40 minutes) the availability of the word’s different meanings is assessed using tasks such as word association or semantic relatedness judgement (Betts et al., 2018; Gaskell et al., 2019; Gilbert et al., 2018; Rodd et al., 2016, 2013). Results from these different paradigms have consistently shown that the availability of the primed, subordinate word meaning is boosted compared to a control, unprimed condition. Importantly, this boost in word-meaning accessibility requires that the target ambiguous wordform is itself present in the prime sentence. For example, the sentence “The man accepted the post in the accountancy firm” will boost the availability of the occupation-related meaning of “post” after 20-40 minutes, whereas in a semantic-priming control condition in which the critical prime word is replaced by a near synonym (e.g., “The man accepted the job in the accountancy firm”), no such boost is observed at this delay (Rodd et al., 2016, 2013). Thus it is participants’ specific experience with the particular ambiguous word that has altered how they are likely to interpret this word when it is subsequently encountered. Word-meaning priming has been observed in carefully controlled lab-

based conditions, as well as in more naturalistic conditions where the primes were presented within short vignettes within a radio programme and priming was measured via a subsequent web-based task (Rodd et al., 2016).

As with the longer-term effects of experience with word meanings that arise incrementally over periods of years, these shorter-term priming effects demonstrate the flexible functionality of the language comprehension system: participants are specifically boosting the availability of those meanings that they are particularly likely to occur (again) in the near future within a particular coherent conversation or text (Rodd et al., 2016). Taken together these findings emphasise the role of prior experience with ambiguous words in lexical-semantic disambiguation in helping comprehenders to rapidly and effectively settle on a single, likely meaning of each word that they encounter. (See Gaskell et al., 2019 for further discussion of likely underlying mechanisms).

(ii) Contextual Cues

The previous section argues for the importance of general, distributional information in guiding word meaning access towards those word meanings that, based on prior experience, are most likely to be encountered in any future context. But such information about the apriori likelihood of any given meaning can, of course, never guarantee accurate comprehension. Comprehenders must make use of additional, context-specific cues to guide them the word meaning that is most compatible with their current situational context. For example, the wordform “bark” should be interpreted very differently when it is preceded by the words “dog” or “tree”.

As with dominance effects, researchers have used a wide range of experimental approaches to assess how and when different types of context guide word-meaning disambiguation (Rodd, 2019). One key theoretical debate has been whether (high-level) information about sentence meaning can directly influence (lower-level) word-meaning access processes. At one end of the theoretical spectrum are ‘exhaustive access’ models (Onifer & Swinney, 1981; Seidenberg, Tanenhaus, Leiman, & Bienkowski, 1982; Swinney, 1979). These models, which have been described as ‘autonomous’ or ‘modular’ (Fodor, 1983) argue for a relatively strict separation between these two aspects of linguistic knowledge, such that all familiar meanings of ambiguous wordforms are automatically accessed, regardless of sentence context. Under this view, context only influences the subsequent meaning selection/integration processes that occur *after* all potential meanings have been accessed. This ‘exhaustive access’ view was supported by early cross-modal semantic priming studies in which participants made responses to visual probe words that followed the ambiguous word (Seidenberg et al., 1982; Swinney, 1979). These studies found that when the probe words were presented immediately at the offset of the ambiguous word then responses were faster for probes that were semantically related to *either* meaning, compared to unrelated probes. This finding is usually interpreted as evidence that both meanings had been automatically activated in response to the ambiguous word. However, subsequent studies found that if the ambiguous word was preceded by a very strongly constraint towards the dominant meaning, then only the contextually appropriate meaning was primed (Tabossi, 1988a; Tabossi & Zardoni, 1993).

This evidence that context *can* modulate the initial access of word meanings is consistent with evidence from eye-tracking studies showing that the time taken to read ambiguous words is strongly modulated by the nature of the preceding context. (See Duffy et al., 2001 for a

comprehensive review of findings concerning how ambiguous words are processed within sentence contexts.) One specific area of debate within the field has concerned the ‘subordinate bias effect’: readers show consistent processing delays when ambiguous words that have a strongly dominant meaning (e.g., “pen”) are preceded by a sentence context that is consistent with an alternative less frequent 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 a limit on how strongly context can modulate word-meaning access: prior context cannot *prevent* readers from accessing strongly preferred, dominant meanings (Kellas & Vu, 1999; Pacht & Rayner, 1993; Rayner, Binder, & Duffy, 1999; Rayner, Pacht, & Duffy, 1994; Sereno, 1995). However, more recently Colbert-Getz and Cook (2013) have shown that a five-sentence strongly constraining context can eliminate the subordinate bias effect. In addition, Leinenger & Rayner (2013) have shown that the subordinate bias effect is reduced when the preceding, biasing context contains the ambiguous word itself.

Taken together, these results confirm the role of sentence context in rapidly guiding readers towards those word meanings that are most likely in a specific sentence context, but that this constraint operates in conjunction with reader’s apriori knowledge about the likelihood of the different meanings. This view is exemplified in the influential Reordered Access Model (Duffy et al., 2001; Duffy, Morris, & Rayner, 1988), which states that (i) whenever an ambiguous word is encountered *all* of its alternative, familiar meanings begin to be automatically activated in parallel, but that (ii) the rate at which these meanings become active is modulated by both sentence context and word-meaning frequency. Under this view, those word meanings that are highly frequent or very strongly consistent with the preceding context become most quickly available for integration into the ongoing discourse. In situations where only one meaning

becomes rapidly available, this meaning can be easily integrated into the higher-level representation of the ongoing discourse without significant disruption to processing (compared with unambiguous words). In contrast, when multiple meanings become available at the same time, then the comprehender must engage in a more time-consuming selection process to determine which meaning is most likely to be intended by the speaker/writer.

More recently, Rodd (2019) has proposed an alternative ‘semantic settling’ framework to describe how the immediate sentence context (and other distributional information) influence word meaning access. This model retains many of the key aspects of the Reordered Access Model, but takes a distributed connectionist perspective in which word-meaning access is viewed as settling process: for each word that is encountered the activation state within a high-dimensional lexical-semantic space settles over time into a stable state that corresponds to a pattern of semantic activation that corresponds to a known, familiar word meaning (Rodd et al., 2004). According to this perspective, multiple cues can influence the settling process by which the system resolves on one of these familiar meanings: these cues include immediate sentence context as well as longer-term distributional knowledge about the likelihoods of different meanings. Critically, this view places learning at its centre: this lexical-semantic space is constantly being re-shaped throughout an individual’s lifespan in response to experience with different word meanings. This re-tuning of lexical semantic knowledge acts to maximize the efficiency of word-meaning access in the individual’s current linguistic environment. This view highlights individual differences in lexical-semantic knowledge: each person’s lexicon is constantly being flexibly re-structured by their specific, idiosyncratic linguistic experiences.

Although these different theoretical perspectives have made clearly stated claims about the mechanisms by which linguistic context can influence word meaning access, these models

are currently somewhat underspecified with respect to claims about exactly what forms of contextual information can influence word meaning access, and how different (potentially conflicting) cues might work together (or in conflict) during word meaning access. Indeed, Rodd (2019) suggests that *any* cue that is a reliable predictor of which meaning is ‘correct’ is likely to be utilized by comprehenders. Under this view there are no structural barriers that can prevent specific (higher-level) forms of knowledge from acting on word-meaning access.

Consistent with this view is the evidence that lexical-semantic access is guided by a wide range of contextual cues. At one end of the spectrum, there is evidence that relatively low-level lexical information about the extent to which word (meanings) tend to co-occur in language cues can guide disambiguation. Witzel & Forster (2015) demonstrated that the interpretation of ambiguous words such as “bat” is strongly influenced by preceding words such as “umpire” that share a strong co-occurrence relationship with the target word (i.e. the baseball related meaning of “bat” tends to co-occur with similar other words to “umpire” because they often occur in similar contexts). Importantly, in this study participants tended to interpret the ambiguous words as being consistent with the meaning that was related to this contextual cue even when the overall sentence context did *not* plausibly support that particular interpretation (e.g., “The umpire tried to swallow the bat but its wings got stuck in his throat”). On the basis of this evidence, Witzel and Forster (2015) propose that word-meaning selection is driven by a “fast-acting, low-level heuristic based on intra-lexical connections”.

However, in addition to this evidence that sentence context guides interpretation via relatively simple information about the different co-occurrence statistics of a word’s different meanings, there is also evidence that higher-level cues about this global interpretation of an ongoing discourse can guide interpretation. For example, Kambe, Rayner & Duffy (2001)

monitored readers' eye movements as they read ambiguous target words that were presented within short paragraph contexts. Critically, they manipulated overall *global* context by using the first sentence of the paragraph to introduce the topic of the paragraph to be consistent with either the dominant or the subordinate meaning the target word ambiguous word. For example, different meanings of the word "band" were introduced by the sentence "Lisa and John spent months looking for the perfect wedding ring" vs. "Lisa and John loved to go to rock concerts with their friends". In addition, they manipulated the *local* context that either preceded or followed the target word, which was always consistent with the subordinate interpretation of the target word (e.g., "It wasn't until they entered Kay's jewelry store that they saw the band."). The results confirm that relatively high-level global contextual information *can* influence disambiguation. Specifically, the global contextual information strongly constrained participants' interpretation of the target ambiguous word when no additional local disambiguating information was available, but had no additional effect when it is consistent with local information. Finally the effect of global context was somewhat delayed effect when it was inconsistent with local information.

In addition to these studies that show that disambiguation can be guided by the local (co-occurrence) or global paragraph/sentence context, recent evidence has indicated that listeners are also influenced by more general background knowledge about the speaker. Specifically, Cai et al., (2017) showed that listeners take account of the accent of the speaker (British versus American) when they are required to disambiguate words that are used somewhat differently by these two dialect communities. Specifically, these studies focused on words like "bonnet" that have different dominant meanings in these two dialects. First, their results (on a range of tasks) confirmed that (consistent with the well attested dominance effects) speakers of British English

can more easily access those word meanings that are dominant in their native dialect (i.e. car part). More importantly, they found that the availability of the alternative, non-preferred subordinate meaning, which is dominant in American English (e.g. type of hat), is increased in British participants, when the words are spoken in an American accent. These data confirm that listener's interpretation are guided not only by the meanings of preceding words, but also by the characteristics of the speaker. The likely mechanisms that drives this effect was further explored in an experiment which compared strongly accented ambiguous words (UK or US) with targets that were presented in a neutral accent. These neutral-accent targets were created by morphing British- and American-accented recordings. These critical items were embedded within contexts of filler words that were all consistently accented with either a UK or US accent. Across a number of experimental conditions they found that the interpretation of these targets was driven by the perceived accent of the filler items and *not* by the specific accent properties of the individual target word. For example, the word "bonnet" was more likely to be given its US meaning when it was surrounded by American-accented speech (compared with a UK accent context), regardless of whether the target word was presented in a US or neutral accent. Cai et al. argue that listeners spontaneously generate an 'accent context' for the speaker that they are currently hearing and use this contextual information to modulate the availability of different word meanings (see Cai et al., 2017 for further details of the likely mechanism by which speaker knowledge can influence lexical processing). They authors also raise the possibility that other cues, beyond accent, may constrain lexical-semantic disambiguation. For example, if other salient information about speakers, such as age or gender, can reliably predict which word meanings are used by individual speakers (e.g. if older and younger participants tend to use the word "tweet" differently) then this information might also contribute to word-meaning access.

Perhaps most intriguing is the (as yet untested) possibility that listeners can retain information about how specific, familiar individuals tend to use specific words (e.g., psychology colleagues may tend to use the statistical meanings of “mean”, “normal” and “significant”). Future work is clearly needed to determine the extent to which listeners (and readers) acquire and maintain such highly specific knowledge about the distributional statistics of how word meanings are used by specific groups or individuals and the extent to which such factors contribute to successful comprehension in more naturalistic paradigms.

Conclusions

Efficient and effective comprehension requires that listeners and readers are able to efficiently perform the mapping from form to meaning. This mapping is made challenging by the ambiguity that is ubiquitous in natural language: accurate processing of the current visual or auditory input alone is insufficient to guarantee that the correct aspects of a word’s meaning are retrieved. A large body of evidence has revealed that word-meaning access is guided by (i) distributional information about the apriori relative likelihoods of different word meanings and (ii) a wide range of contextual cues that indicate which meanings are most likely in the current context. Successful comprehension is the result of the ability to rapidly and fluently integrate these cues for every word that is encountered such that only relevant aspects of word meaning are integrated into the ongoing representation of the overall meaning of the discourse.

References

Armstrong, B. C., & Plaut, D. C. (2008). Settling dynamics in distributed Nntworks explain yask 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>
- Betts, H. N., Gilbert, R. A., Cai, Z. G., Okedara, Z. B., & Rodd, J. M. (2018). 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>
- Carston, R. (2019). Ad Hoc Concepts, Polysemy and the Lexicon. In B. Clark, K. Scott, & R. Carston (Eds.), *Relevance, Pragmatics and Interpretation* (pp. 150–162). <https://doi.org/DOI: 10.1017/9781108290593.014>
- Castles, A., Rastle, K., & Nation, K. (2018). Ending the Reading Wars: Reading Acquisition From Novice to Expert. *Psychological Science in the Public Interest*, 19(1), 5–51. <https://doi.org/10.1177/1529100618772271>
- Colbert-Getz, J., & Cook, A. E. (2013). Revisiting effects of contextual strength on the subordinate bias effect: Evidence from eye movements. *Memory and Cognition*, 41(8),

1172–1184. <https://doi.org/10.3758/s13421-013-0328-3>

- Coltheart, M., Rastle, K., Perry, C., Langdon, R., & Ziegler, J. (2001). DRC: A dual route cascaded model of visual word recognition and reading aloud. *Psychological Review*, *108*(1).
- 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>
- 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). <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)
- Ferreira, F., Bailey, K. G. D., & Ferraro, V. (2002). Good-enough representations in language comprehension. *Current Directions in Psychological Science*, *11*(1). <https://doi.org/10.1111/1467-8721.00158>
- Fodor, J. A. (1983). *The Modularity of Mind*. Cambridge, MA: MIT Press.
- 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>
- 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*, *44*(10), 1533–1561.

<https://doi.org/10.1037/xlm0000532>

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)

Johnsrude, I. S., & Rodd, J. M. (2016). Factors That Increase Processing Demands When

Listening to Speech. In G. Hickok & S. L. Small (Eds.), *Neurobiology of Language* (pp.

491–502). <https://doi.org/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>

Leinenger, M., & Rayner, K. (2013). Eye movements while reading biased homographs: Effects

of prior encounter and biasing context on reducing the subordinate bias effect. *Journal of*

Cognitive Psychology, 25(6), 665–681. <https://doi.org/10.1080/20445911.2013.806513>

- Maciejewski, G., Rodd, J. M., Mon-Williams, M., & Klepousniotou, E. (2018). The cost of learning new meanings for familiar words. Retrieved from <https://doi.org/10.31234/osf.io/7ydkw>
- Marslen-Wilson, W. D., & Welsh, A. (1978). Processing interactions and lexical access during word recognition in continuous speech. *Cognitive Psychology*, 10(1), 29–63. [https://doi.org/10.1016/0010-0285\(78\)90018-X](https://doi.org/10.1016/0010-0285(78)90018-X)
- McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: I. An account of basic findings. *Psychological Review*, 88(5). <https://doi.org/10.1037/0033-295X.88.5.375>
- Mirković, J., & Altmann, G. T. M. (2019). Unfolding meaning in context: The dynamics of conceptual similarity. *Cognition*, 183, 19–43. <https://doi.org/10.1016/j.cognition.2018.10.018>
- Morton, J. (1969). Interaction of information in word recognition. *Psychological Review*, 76(2), 165–178. <https://doi.org/10.1037/h0027366>
- 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).
- Ramiro, C., Srinivasan, M., Malt, B. C., & Xu, Y. (2018). Algorithms in the historical emergence of word senses. *Proceedings of the National Academy of Sciences of the United States of America*, 115(10), 2323–2328. <https://doi.org/10.1073/pnas.1714730115>
- Rayner, K. (1998). Eye Movements in Reading and Information Processing: 20 Years of Research. *Psychological Bulletin*, 124(3).
- 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. (2019). *Settling into Semantic Space: An Ambiguity-Focused Account of Word-Meaning Access*. <https://doi.org/10.31234/OSF.IO/75M4K>
- 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), 1095–1108. <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, 16–37. <https://doi.org/10.1016/j.jml.2015.10.006>

- Rodd, J. M., Gaskell, M. G., & Marslen-Wilson, W. (2002). Making sense of semantic ambiguity: Semantic competition in lexical access. *Journal of Memory and Language*, 46(2), 245–266. <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), 89–104. <https://doi.org/10.1016/j.cogsci.2003.08.002>
- 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), 180–198. <https://doi.org/10.1016/j.jml.2012.08.002>
- Seidenberg, M. S., & McClelland, J. L. (1989). A Distributed, Developmental Model of Word Recognition and Naming. *Psychological Review*, 96(4).
- 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>
- 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>
- 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)

Tabossi, P. (1988a). 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. (1988b). Effects of Context on the Immediate Interpretation of Unambiguous Nouns.

Journal of Experimental Psychology: Learning, Memory, and Cognition, 14(1), 153–162.
<https://doi.org/10.1037/0278-7393.14.1.153>

Tabossi, P., & Zardon, F. (1993). Processing Ambiguous Words in Context. *Journal of Memory and Language*, 32(3). <https://doi.org/10.1006/jmla.1993.1019>

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., & Rodd, J. M. (2015). Resolving Semantic Ambiguities in Sentences: Cognitive Processes and Brain Mechanisms. *Linguistics and Language Compass*, 9(10), 391–405.

<https://doi.org/10.1111/lnc3.12160>

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>

Yee, E., & Thompson-Schill, S. L. (2016). Putting concepts into context. *Psychonomic Bulletin and Review*, 23(4), 1015–1027. <https://doi.org/10.3758/s13423-015-0948-7>