Published in final edited form as:

Brain Res. 2010 December 13; 1365: 66–81. doi:10.1016/j.brainres.2010.09.057.

Are There Mental Lexicons? The Role of Semantics in Lexical Decision

Katia Dilkina¹, James L. McClelland², and David C. Plaut³

¹ University of Victoria, Victoria, BC, Canada

² Stanford University, Stanford, CA, USA

³ Carnegie Mellon University, Pittsburgh, PA, USA

Abstract

What is the underlying representation of lexical knowledge? How do we know whether a given string of letters is a word, whereas another string of letters is not? There are two competing models of lexical processing in the literature. The first proposes that we rely on mental lexicons. The second claims there are no mental lexicons; we identify certain items as words based on semantic knowledge. Thus, the former approach – the multiple-systems view – posits that lexical and semantic processing are subserved by separate systems, whereas the latter approach – the single-system view – holds that the two are interdependent.

Semantic dementia patients, who have a cross-modal semantic impairment, show an accompanying and related lexical deficit. These findings support the single-system approach. However, a report of an SD patient whose impairment on lexical decision was not related to his semantic deficits in item-specific ways has presented a challenge to this view. If the two types of processing rely on a common system, then shouldn't damage impair the same items on all tasks?

We present a single-system model of lexical and semantic processing, where there are no lexicons, and performance on lexical decision involves the activation of semantic representations. We show how, when these representations are damaged, accuracy on semantic and lexical tasks falls off together, but not necessarily on the same set of items. These findings are congruent with the patient data. We provide an explicit explanation of this pattern of results in our model, by defining and measuring the effects of two orthogonal factors – spelling consistency and concept consistency.

Keywords

semantics; lexicon; lexical decision; word reading; computational modeling; semantic dementia

1. Introduction

The long-standing tradition in psycholinguistics has been to explain linguistic processes in terms of a dual system consisting of a powerful rule-based component, which handles the

Correspondence regarding this manuscript should be directed to Katia Dilkina, Cornett Building A234, 3800 Finnerty Road, Victoria, BC, V8P 5C2, Canada. Phone number: 250-853-3762. katia.dilkina@gmail.com.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

regular cases, along with a repository which holds specific entries and manages exceptional items (i.e., items that do not follow the rules). A classic task thought to rely on one such repository is lexical decision, where the participant is presented with letter strings that are either real words or word-like nonwords. In the yes/no version of the task, the subject has to verify whether each item is a word or not. In the two-alternative forced-choice version (2AFC), words and orthographically-similar nonwords are presented in pairs, and the subject is asked to select the real word.

According to theories that postulate the existence of mental lexicons (e.g., Caramazza, 1997; Coltheart et al., 2001; Levelt, 1989), this task is performed by accessing the orthographic input lexicon and checking for the appropriate entry. If the given string exists in the mental lexicon, then it is a word; if it does not, then it is a nonword. If the lexicon is damaged then performance will be impaired.

The alternative view is that there are no word-form representations per se, and in order to evaluate whether a string is a word or not, one needs to access the semantic system (Dell & O Seaghdha, 1992; Patterson et al., 2006; Plaut, 1995, 1997; Plaut et al., 1996; Rogers et al., 2004b). This approach has been supported by numerous studies showing semantic effects including priming, concreteness/imageability, and valence, in lexical decision (e.g., Azuma & Van Orden, 1997; Binder et al., 2003; Kuchinke et al., 2005; Pexman, Hargreaves, Edwards, Henry, & Goodyear, 2007; Samson & Pillon, 2004). Importantly, at least some of those studies have emphasized that semantic factors are at play *concurrently* with word identification and not simply accessed after lexical processing (Forster & Hector, 2002; Marslen-Wilson, 1987; Tyler, Moss, Galpin, & Voice, 2002; Wurm, Vakoch, & Seaman, 2004).

Strong evidence in favor of the notion that (a) lexical and semantic knowledge rely on a common cognitive system, and (b) regular and exception items are not processed via separate routes, comes from individuals with semantic dementia. Semantic dementia (SD) is characterized by progressive atrophy of the anterior temporal cortex accompanied by increasing deficits of conceptual knowledge (Knibb, Kipps, & Hodges, 2006; Neary et al., 1998). SD patients exhibit a particular profile of conceptual deficits accompanied by a similar, and related, profile of lexical deficits. General aspects of meaning remain relatively preserved while specific details are lost (e.g., Hodges, Graham, & Patterson, 1995; Schwartz, Marin, & Saffran, 1979). Also, frequently occurring and/or typical items and features deteriorate more slowly than rare and/or atypical ones but eventually both types of items do deteriorate (e.g. Hodges et al., 1995; Lambon Ralph, Graham, & Patterson, 1999; Papagno & Capitani, 2001). Similarly, in lexical decision, the majority of SD patients show a significant deficit which correlates with the degree of semantic impairment, as indicated by scores on naming, word-picture matching, and the Pyramids & Palm Trees test of associative semantic knowledge (Benedet, Patterson, Gomez-Pastor, & de la Rocha, 2006; Funnell, 1996; Graham et al., 2000; Knott et al., 1997, 2000; Lambon Ralph & Howard, 2000; Moss & Tyler, 1995; Patterson et al., 2006; Rochon et al., 2004; Rogers et al., 2004b; Saffran et al., 2003; Tyler & Moss, 1998). Lexical decision performance is strongly influenced by word frequency (Blazely et al., 2005; Rogers et al., 2004b). In the yes/no paradigm, the responses are often not random; instead, they are biased towards 'yes' on both words and nonwords, thus yielding a large number of false positives (Funnell, 1996; Knott et al., 1997; Saffran et al., 2003; Tyler & Moss, 1998). In the 2AFC paradigm, the patients show a marked typicality effect such that when the spelling of the word is more typical than the spelling of the accompanying nonword, they tend to perform well; but in the reverse case, they show a deficit (Patterson et al., 2006; Rogers et al., 2004b). In fact, the individuals with the most severe semantic deterioration reliably prefer the nonword to the actual word in such pairs (Rogers et al., 2004b).

In summary, SD patients' performance on lexical decision (1) usually falls below the normal range, and declines over the course of the disease; (2) is correlated with the degree of semantic impairment; (3) shows considerable frequency and typicality effects (which is seen in these patients' performance on all tasks that rely on semantic and/or lexical knowledge). All of these results taken together strongly suggest that lexical decision is not divorced from semantic processing, and intact conceptual representations are necessary for adequate performance. These findings from semantic dementia undermine the necessity for – and even the plausibility of – lexical-level representations independent of semantic representations.

Still, this is not the end of this debate. Neuropsychological data inevitably exhibits substantial variability, opening room for alternative interpretations. In a study of two semantic dementia patients, Blazely, Coltheart, and Casey (2005) presented two challenges to the single-system account of lexical and semantic processing. The two patients reported, EM and PC, appeared to have comparable semantic deficits as assessed by picture naming, spoken and written word-picture matching, and name comprehension. However, while patient PC was also impaired on two non-semantic tasks, word reading and lexical decision, patient EM was not. Thus, the first challenge to the single-system model was based on the finding of dissociation between semantic performance and lexical performance. In addition, patient PC's poor performance on lexical decision did not show significant item-by-item correlation with any semantic task but written word-picture matching (cross-test analyses were not reported for word reading). The authors argued that if lexical and semantic processing are subserved by a common system then (1) when semantics is impaired, lexical performance should also be impaired (thus, there should be no cases of dissociation such as patient EM), and (2) impairment across semantic and lexical tasks should affect an overlapping set of items. The fact that patient PC made lexical decision errors on items different from those which he failed on semantically was taken to suggest that the two deficits are unrelated; they are due to damage to two distinct systems rather than a single system.

We and our colleagues have already addressed the argument that dissociation between lexical and semantic performance entails separate systems (Dilkina, McClelland, & Plaut, 2008; Woollams, Lambon Ralph, Plaut, & Patterson, 2007) by demonstrating that cases of apparent dissociation can arise from individual differences within a single system. People vary in their abilities, skills, experience, and even brain morphology. Brain-damaged patients further vary in the extent and anatomical distribution of their atrophy as well as in their post-morbid efforts to sustain affected domains of performance. These differences can lead to apparent dissociations, even within a single-system account. This statement has been supported by simulation experiments in the triangle model of word reading (Plaut, 1997; Woollams et al., 2007) as well as in a recent extension of this approach which implements semantics (Dilkina et al., 2008). These investigations have shown that individual differences can meaningfully and successfully explain the findings of both association and dissociation between semantic impairment (as measured by picture naming and word-picture matching) and lexical impairment (as measured by word reading, thought to depend on the lexicon in dual-system approaches) in SD patients. Now, we will tackle these issues as they arise in the lexical decision task, focusing on the pattern of correlation and item consistency between performance on lexical decision and other tasks.

How can the single-system view explain these two seemingly contradictory sets of results – a significant **overall correlation** between accuracy on lexical decision with atypical items and accuracy on semantic tasks (e.g., Patterson et al., 2006) but no **item-specific correlation** between lexical decision and the same semantic tasks (Blazely et al., 2005)? These findings can be accounted for in the following way: Within the proposed single

system, word recognition in lexical decision is indeed affected by semantic damage. Words with irregular pronunciation and/or inconsistent spelling are affected more because they are incoherent with the general grapheme and phoneme co-occurrence statistics of the language and identifying each of them as a familiar sequence of letters or sounds relies more heavily on the successful activation of semantic knowledge corresponding to this sequence of letters or sounds. When the semantic representations are degraded, the system is unable to identify such strings as words. Our theory then makes two predictions: (1) there should be a significant overall correlation between accuracy in semantic tasks and accuracy in lexical decision (as reported by Patterson et al., 2006); (2) there should be no significant correlation between items impaired in semantic tasks and items impaired in lexical decision because performance levels are governed by two orthogonal dimensions, which we formally define later in the paper – *spelling consistency* in lexical decision and *concept consistency* in semantic tasks. However, this in-principle argument remains to be demonstrated empirically within our single-system model. In the current investigation, we confirmed both of these hypotheses.

It should also be noted that our theory does predict item-specific correlations between reading and spelling atypical items and the semantic integrity of these items, in particular when tested on tasks such as naming and word-picture matching. Why does it make this prediction? How are these lexical tasks different from lexical decision? The main distinction is that reading and spelling require not merely the activation of a discernible semantic representation but also the consequent activation of another surface form. As the success of this mapping depends on the integrity of the semantic representation, items that rely on the semantic route for reading and spelling – which is all words but especially those with atypical spelling-to-sound/sound-to-spelling correspondence (called irregular and inconsistent, respectively, as defined in the Experimental procedure section) – would show item-by-item correlation between semantic and lexical tasks. This correlation would be particularly prominent for semantic tasks that include words as input or output as such tasks partially share pathways with reading and spelling. Again, this is only a theoretical argument that needs to be demonstrated in the actual single-system model. We have done so in the present investigation.

In summary, this project investigates the relationship between semantic and lexical processing in a model where knowledge of both types of items is stored in a single integrated system. In the model we describe, we first explore how lexical decision is performed, without a lexicon; then go on to address how lexical decision performance relates to performance on other tasks.

2. Results

The neural network (Figure 1) and the materials used for the present simulations are described in detail in the Experimental procedure section at the end of the paper. The most important aspect of the architecture was that the different types of surface representations connected to a common cross-modal hub, dubbed the integrative layer, which we view as functionally analogous to the anterior temporal cortex in the human brain – the area first and most severely affected in patients with semantic dementia. In addition, the orthographic and

¹Throughout this article, we use the terms ‘regular’ and ‘irregular’ to refer to the pronunciation of a written word, even though, in our view, there are no rules that define what is regular or not; we use this terminology because it is frequently used in other studies investigating word reading, where performance on lists of ‘regular’ and ‘irregular’ words are compared. In fact it is our view that the true underlying factor affecting reading performance is not conformity to a specific set of rules but consistency in the mapping from spelling to sound (see Plaut et al., 1996). Notably, there are two distinct types of consistency in the relationship between orthography and phonology – spelling-to-sound consistency (which here we call ‘regularity’) and sound-to-spelling consistency (which we call ‘spelling consistency’ or, when the context is clear, ‘consistency’ for brevity).

phonological surface representations have a direct link via a smaller hidden layer (functionally analogous to the left posterior superior temporal/angular region known to be involved in reading of both words and pronounceable nonwords). The model was tested on adapted versions of the following tasks commonly used with SD patients: picture naming, spoken word-picture matching, Pyramids and Palm Trees with pictures and with words, single word reading, single word spelling, and 2AFClexical decision with pseudohomophones. In our simulations, all of these tasks required the involvement of the integrative layer.

The focus of this investigation was the lexical decision task, in which the network was presented with pairs of items – a word and a nonword that differed in spelling but had a common pronunciation. The stimuli were presented one at a time at the orthographic layer; activation spread through the system and – after the input was removed – the network re-activated the representation at orthography. We used the relative strength of this “orthographic echo” as a measure of lexical decision.

Baseline performance & lexical decision in our single-system model

At the end of training, the network had an accuracy of 100% on all tasks, including lexical decision. Even though the model had no explicit word-level representations, it was perfectly able to distinguish between real words and nonwords. Furthermore, the lexical decision task used was very difficult in that the word-nonword pairs did not differ phonologically (i.e. they were homophonic) and differed orthographically in only one letter. Moreover, in about half of the items, the nonword had a more consistent spelling than the word ($W < NW$). Still, the network was able to discern the real word. It achieved this by spreading activation from the orthographic layer to semantics and the direct layer, and then back to orthography.

Within our framework, where there is no list of lexical entries, even though lexical decision is not a *purely* semantic task, and on some occasion it can be performed without semantic access, it is a *necessarily* semantic task – in order to perform accurately on all types of trials, there needs to be semantic access. Words with irregular pronunciation and/or inconsistent spelling are particularly reliant on semantic activation, and therefore most vulnerable to semantic damage (see also Jefferies, Lambon Ralph, Jones, Bateman, & Patterson, 2004; Patterson et al., 2006; Plaut, 1997; Rogers et al., 2004b).

In order to evaluate the relative contributions of the semantic pathway and the direct pathway to lexical decision performance, we collected data when either only the direct layer, or only the integrative layer, or both participated during testing. The results can be seen in Figure 2. The first thing to note is that at the end of training, prior to damage, the integrative layer was able to perform lexical decision perfectly on its own but the direct pathway was not. The direct pathway does not have full proficiency on either $W > NW$ or $W < NW$ trials, and it does show differential performance on the two types of trials even before the network is damaged.

Why isn't the lexical competence of the direct pathway flawless at baseline? One reason is because it is predominantly trained in conjunction with the integrative layer. At the beginning of the second stage of training (which is when the direct pathway becomes involved), the integrative layer already “knows” the words in the model's vocabulary because the pathway between it and phonology has already been trained. As training proceeds, on reading and spelling trials the semantic pathway (i.e. between orthography and phonology through the integrative layer) is already able to contribute something because of the available connections between phonology and the integrative layer, as well as the developing connections with orthography which are acquired on trials other than reading and spelling (e.g., when mapping from orthography to vision; reading a word and imagining

its referent). The direct pathway, on the other hand, does not yet have any lexical knowledge. As training proceeds, both pathways learn in conjunction. Because of the generally systematic relation between spelling and pronunciation, the direct pathway – linking orthography and phonology – becomes very sensitive to regularity and spelling consistency. So even though there are trials on which the direct pathway is trained alone (one third of all lexical trials), they serve mostly the regular and/or consistent items. Thus, at the end of training, performance on lexical tasks – lexical decision as well as word reading and spelling – is at ceiling when both the integrative and the direct layer participate or when the integrative layer participates alone, but **not** when the direct pathway is on its own.

Furthermore, even if the direct pathway were at ceiling at the end of training, this does not mean that when the integrative layer is damaged, lexical performance in the direct pathway will remain intact. This is because *damage at semantics does not simply diminish the model's overall lexical competency; it introduces noise which interferes with the competency of the direct pathway* (see also Woollams et al., 2007). This is an aspect of the model that distinguishes our approach from the multiple-systems approach. In our model, the semantic and direct routes used to perform lexical tasks are not encapsulated modules. They develop together and therefore become interdependent; damaging one affects the performance of the other.

As can be seen in Figure 2, the progressive semantic lesion affected **both** pathways but in different ways. The integrative layer continued to perform well on W>NW pairs, even though eventually they did suffer some impairment. In contrast, W<NW pairs started declining very quickly and continued to fall off past chance level. The direct layer, on the other hand, was initially able to sustain its performance on both types of items at a relatively constant level. At more advanced stages of semantic damage, however, W<NW showed a slight impairment, while W>NW actually showed improvement. This is not surprising because as the integrative layer became increasingly limited in its resources, it fell on the direct route to handle as much of the lexical tasks as possible. However, because its ability is also not perfect, the further retraining during damage improved mostly the words with regular pronunciation and consistent spelling.

Because of these trends, the difference in accuracy between W>NW and W<NW trials was strongest when only the integrative layer participated in lexical decision, which can be seen in Figure 3 (main effect of spelling consistency: W>NW > W<NW, $F(1,19)=1060.8$, $p<.0001$; main effect of testing method: Integrative & Direct > Integrative only > Direct only, $F(1,19)=913.09$, $p<.0001$; interaction: $F(2,38)=188.17$, $p<.0001$).

In summary, our exploration of the two pathways underlying lexical decision performance in our model revealed that even when damage occurs only in the semantic system, the direct pathway is also affected. Furthermore, the retraining regime used between lesions affected overall lexical decision accuracy as well as the effect of spelling consistency in lexical decision in unique ways in the two pathways. By virtue of the general processing mechanisms working in connectionist models, damage in the integrative layer, postmorbid retraining, and decay in the direct layer combine to produce a strong preference for letter strings with consistent spelling over those with inconsistent spelling. Thus, even though our model has a semantic as well as a direct route for lexical tasks, these are by no means separate systems. They are highly interactive and together change in the face of damage.

Overall & item-specific correlations across tasks

In agreement with SD patient reports (e.g., Patterson et al., 2006), our overall correlational analysis showed that the model's accuracy on lexical decision with pairs where the word had a less consistent spelling than the nonword (W<NW) correlated strongly not only with the

accuracy on the other two lexical tasks using atypical words, reading and spelling ($p < .0001$; Figure 4), but also with the accuracy seen on all of the semantic tasks (all $p < .0001$; Figure 4).

In order to further evaluate the relationship among semantic and lexical tasks in our model, we performed a series of item-by-item correlational analyses to see whether damage to the integrative layer caused the same items to be impaired on semantic tasks on one hand and lexical tasks on the other. Like SD patient s (e.g., Bozeat et al., 2000; Hodges et al., 1995), our model's impairment on all semantic tasks included an overlapping set of items (all $p < .005$; Table 1). Furthermore, in agreement with the patient data reported by Graham et al. (1994) and Jefferies et al. (2004), the model's reading performance showed item-specific correlations with its performance on WPM, picture naming, PPT with words (all $p < .05$), and marginally so with PPT with pictures ($p = .07$). In addition, we found significant item-specific correlations between spelling and all four of the semantic tasks (all $p < .05$). Not surprisingly, we also found that reading and spelling exhibited high item consistency ($p < .0001$).

Most importantly, in line with Blazely et al.'s (2005) report of patient PC, and contrary to what the authors thought our single system account would predict, we found **no** relationship between the items impaired on any of the semantic tasks and the items impaired on lexical decision (for either only the W<NW trials, or all trials; see Table 1). The only task showing item-by-item correlation with lexical decision was spelling.

So why is it that although overall lexical decision performance correlated positively with overall performance on all other tasks in our model (all $p < .0001$; Figure 4), it only exhibited an item-specific relationship with word spelling? In order to understand this, one needs to understand the learning mechanism that operates in connectionist networks and the factors these mechanisms are sensitive to. Because the integrative layer is used to mediate among all surface forms, both object-related (visual, action/encyclopedic) and word-related (orthographic, phonological), the learned semantic representations are sensitive to any and all partially or fully consistent relationships present in the mappings among these surface representations. The mappings between object-related surface representations and word-related surface representations are arbitrary, so there is little or no consistency there. There **is** some degree of consistency, however, in the mappings between orthography and phonology, and between the visual features and the action/encyclopedic features. In what follows we consider measures of these two kinds of consistency.

We make use of a measure of the consistency in mapping from phonology to orthography (see Table 2 for a distribution of consistent vs. inconsistent lexical items), which we call *spelling consistency*. The term can be easily understood as the answer to the question: Given a certain spoken word, what is its written form? In English (and, by design, in our model), more often than not, there are multiple possible spellings for a particular pronunciation. Only one of those is considered *consistent*, and that is the written form that occurs most often for words that have the same phonological rime (vowel plus subsequent consonants) as the spoken word in question. The rest of the candidates are considered inconsistent. In reality, this dimension is continuous rather dichotomous but we are using a simplified classification for the purposes of this analysis.

The notion of spelling consistency should not be confused with the notions of regularity and consistency as they apply to the mapping from spelling to sound, as explored at length in earlier publications (Plaut et al., 1996). Spelling consistency refers to the status of a possible spelling given a particular pronunciation (phonology to orthography), whereas regularity and spelling-to-sound consistency refer to the status of a possible pronunciation given a particular spelling (orthography to phonology). See later discussion with respect to reading

as well as the description of the materials in the Experimental procedure section. The notion is also different from, but is likely correlated with, the concept of orthographic typicality used by Rogers and colleagues (2004b).

The second measure we use, *concept consistency*, is defined with respect to the relationship between an object's visual properties and its other semantic (action and encyclopedic) properties. We first consider the extent to which each item is typical or atypical relative to its category in its visual and other semantic features. We measure this by calculating the average visual and average action/encyclopedic vector for each category, then taking the cosine of each specific item's vector and its category vector. Items with a cosine value higher than the average are categorized as typical of their category, while those with lower value are categorized as atypical. Once again, these are continuous dimensions which have been dichotomized for simplicity. Now, to define each concept as consistent or inconsistent, we need to consider both the visual and the other semantic typicality values. Items are treated as *semantically consistent* if they are either typical in both their visual and other semantics or if they are atypical in both.

The measure of semantic consistency just defined is different from other measures that have been used in previous literature. Most often, researchers speak of the typicality of objects with respect to the superordinate categories to which they belong (Rosch et al., 1976), and/or consider individual visual features of an object that may be typical or atypical of their superordinate category (Rogers et al., 2004b). A measure based solely on visual typicality may be especially useful for capturing performance in the visual object decision task used by Rogers et al. (2004b). We have found, however, that in the range of tasks considered here, visually atypical items that are also atypical in their other semantic properties benefit from this double atypicality, supporting the choice of measure we use here.

Having defined spelling consistency and concept consistency, we can now ask whether and how each of these dimensions affects performance on our tasks. We used these constructs as parameters in an analysis of variance of the simulation data for each of the seven tasks. The results are presented in Figure 5. Not surprisingly, spelling consistency had a significant positive effect on all lexical tasks as well as on PPT with words – consistent words were more robust to damage than inconsistent ones. The effect was strongest in word spelling, followed by lexical decision. There was no significant effect of spelling consistency in either WPM or PPT with pictures, and in naming the effect was significant but went in the opposite direction! Concept consistency, on the other hand, affected positively all tasks including lexical decision. Importantly, there was a significant interaction between the two dimensions in all tasks **but** lexical decision. The interaction was such that the concept consistency effect was much stronger for the items whose spelling was inconsistent than those whose spelling was consistent (Figure 5). This can be explained by the fact that atypical words (irregular and/or inconsistent) rely more on the semantic route in reading and spelling. Therefore, they have become more sensitive to factors governing the learned representations in the integrative layer, and these are captured by concept consistency. Similar logic may explain why spelling consistency had a negative main effect for naming, which was the most sensitive semantic task: items with inconsistent spelling, which most heavily rely on the semantic route, performed better in naming than items with consistent spelling. Perhaps this occurred because the error -correcting learning mechanism in the model strengthened these items' semantic representations to deal with them effectively in lexical tasks.

These findings show that the two orthogonal dimensions that govern the nature of the representations in the integrative layer affect performance on the various tasks in different ways. Semantic tasks are influenced strongly and positively by consistency in the mapping

between the visual features and action/encyclopedic features. They are generally insensitive to spelling consistency, even though the results for naming – the only non-multiple-choice semantic task – suggested that items with inconsistent spelling may in fact have an advantage in semantic tasks. Of course, this is only true when the items are not presented as written words (or required as written output). In that case, as in PPT with words, there is preference for the words with consistent spelling. All lexical tasks also exhibit a strong advantage of consistent words. In addition, reading and spelling are also affected positively by concept consistency, especially for items with inconsistent spelling, which was also the case in all semantic tasks.

Our most important finding was that lexical decision was only weakly sensitive to concept consistency, as evidenced by the small main effect of this variable (compared to the other two lexical tasks) and the lack of interaction between it and spelling consistency. The fact that semantic tasks are strongly governed by concept consistency and indifferent or negatively affected by spelling consistency, while lexical decision is strongly positively affected by spelling consistency and only weakly by concept consistency explains why lexical decision did not show an item-specific relationship with any of the semantic tasks.²

One might wonder why we did not find an item-specific relationship between lexical decision and reading but we did find one between lexical decision and spelling. The reason is that while both lexical decision and spelling are strongly affected by spelling consistency (that is, how typical is the spelling of a word given its pronunciation), reading is more strongly affected by *regularity* (how typical is the pronunciation of a word given its spelling). As discussed earlier, words can be regular but inconsistent, for example ‘byte’. When we analyzed our simulated reading data using regularity instead of spelling consistency, we found a much stronger effect as well as a much more robust interaction with concept consistency (see Figure 5).

Finally, there is the issue of how semantic performance relates to performance on lexical decision versus word reading. In our model, the former task shows an overall correlation but not an item-specific correlation with semantic tasks whereas the latter task shows both. We explained these findings in our model in terms of sensitivity to spelling consistency and concept consistency. What about patient performance?

Even though Blazely and colleagues (2005) tested their patients on both tasks, they only reported item-by-item correlations for lexical decision. However, in an earlier study of three SD patients tested on the very same tasks, Graham and colleagues (1994) found significant item-specific correlations between the patients’ reading of irregular words and their performance on word-picture matching as well as name comprehension. They found no such relationship between reading regular words and semantic performance. More recently, in a thorough investigation of reading in seven SD patients, McKay et al. (2007) looked at both accuracy and response latency. Using accuracy, they found an item-specific relationship

²Notably, Blazely et al. (2005) did find an item-by-item correlation between lexical decision and one semantic task: WPM with written words. The investigators attributed this relationship to the fact that both tasks use written words as input. While our set of simulations didn’t include a written version of WPM, we did have another semantic task that used written words as input – PPT with words. This task did show a positive effect of spelling consistency. Why did it not show item-by-item correlation with lexical decision? There are at least three reasons: Firstly, the effect of spelling consistency was not nearly as strong in PPTw as in lexical decision. Secondly, PPTw was also affected by concept consistency, and it showed a strong interaction between the two variables, while lexical decision did not. Finally, as explained in the Experimental procedure section, our model was tested on multiple-choice semantic tasks (PPT and WPM) with the participation of only the integrative and the action/encyclopedic layers. This testing procedure ensured that a common set of resources is available across the different trials and inputs, but unfortunately it minimized the effects of testing modality. Thus, while we believe that there is *some* relationship between PPTw and lexical decision by virtue of a shared pathway in processing (O→S), we did not find this relationship to be significant at the item-specific level due to the reasons outlined above.

between successfully reading irregular words and knowing their referents, which was assessed with two tasks: a free-response concept definition task and a multiple-choice task where the patients were provided with a definition and had to choose the corresponding concept among three words: target, semantically-related foil and semantically-unrelated foil. Similarly to Graham et al. (1994), the researchers found no such relationship with respect to regular words. However, when they looked at reading latencies – a much more sensitive measure than accuracy – they found item-by-item correlations between reading and semantic impairment on both regular and irregular words. These results are very much in line with our view that lexical tasks in general rely on the semantic system and that irregular/inconsistent words are particularly dependent on semantic knowledge.

In contrast, according to the multiple-systems framework, the often-seen lexical deficits in semantic dementia arise not because of the semantic damage but because of damage to the neighboring lexical system. As a result, items impaired on semantic tasks and items impaired on lexical tasks should be unrelated.³ This kind of reasoning applies not only to lexical decision but also to reading. Thus, the fact that the patients' reading performance does seem to be related to the semantic performance in item-specific ways is problematic for this view. This pattern of results, however, is expected by the single-system view, and was indeed present in our model.

3. Discussion

The purpose of this investigation was twofold. First, we wanted to demonstrate how a system with no lexicons can perform a difficult lexical decision task involving fine discriminations between words and nonwords. The single-system model has been successfully tested on lexical decision tasks before (e.g., Plaut, 1997) but we wanted to implement a version of the task that most closely corresponded to the one used with SD patients. The second goal was to explain why, within our single-system framework, patients' lexical decision accuracy correlates with their semantic accuracy, but there need not be an item-specific relationship between performance in lexical and semantic tasks.

The results showed that our single-system model was indeed able to successfully perform the two-alternative forced choice lexical decision task in the absence of word-level representations. To do so, it relied both on semantic knowledge and on knowledge of the mappings between orthography and phonology. We evaluated the relative contribution of the shared integrative layer versus the direct layer between orthography and phonology, and found that – prior to damage – the integrative layer can perform the task perfectly on its own, while the direct layer cannot. This is due to the fact that the two pathways are interdependent, and even though the direct layer is trained to perform lexical mappings on its own, most of the time it does so in conjunction with the integrative layer. We also found that both layers are sensitive to differences in spelling consistency, and that sensitivity increases with damage. In the integrative pathway, this increase is due to the quick and sharp decline in performance on trials with $W < NW$, which actually fall below chance level (also seen in SD patient data; cf. Rogers et al., 2004b). In the direct pathway, this increase is due to the improvement on trials with $W > NW$. These findings illustrate that (1) lexicons are not necessary for lexical decision, no matter how difficult the task; (2) within our single-system model, the shared cross-modal integrative layer and the direct layer connecting orthography and phonology contribute to lexical decision performance in distinct ways; (3) the two routes subserving word processing are highly interactive.

³Notably, since performance on both types of tasks is sensitive to frequency, an association may arise due to frequency effects. This is why it is important to control for this variable in overall correlational analyses and/or look at performance across tasks at the item-specific level.

A recent imaging study using a combination of fMRI and MEG to localize semantic activation in the brain during a yes/no lexical decision task (Fujimaki et al., 2009) lends further support to the notion that this task engages the same system as semantic tasks, namely cross-modal conceptual representations in the anterior temporal lobes. Fujimaki and colleagues (2009) tested their participants on lexical decision as well as phonological decision – detecting the presence of a vowel in a visually presented character. They reasoned that both tasks require orthographic and phonological processing, but only lexical decision involves also semantic activation. They found that in the time window of 200–250 msec after stimulus onset, the only area that showed a significant difference in activation between the two tasks was the anterior temporal lobe. As mentioned earlier, this is indeed the brain region consistently and most severely compromised in semantic dementia. Activation in this area was greater for lexical decision compared to phonological decision. Follow-up comparisons indicated that this difference remained significant between 200msec and 400msec after stimulus onset. The authors concluded that there is semantic access during lexical decision and it occurs in an early time window, as early as 200msec after the presentation of the written word.

Similar results were found in an ERP study, which adopted a different approach. Hauk et al. (2006) used a yes/no lexical decision paradigm with stimuli identical to the ones used with SD patients in the past (cf. Rogers et al., 2004b) – namely, pairs of words and homophonic nonwords. The investigators looked at the time course and localization of the effects of spelling typicality versus lexicality during task performance, and found that the two variables interacted at around 200 msec in that there was a strong effect of typicality only for the real words; this interaction was localized in the left anterior temporal area. These results provide further support not only to the idea that lexical decision involves semantic activation and the localization of this activation to the left anterior temporal cortex, but also to the notion that it is the atypical words in particular that most heavily require such activation – which is what we observed in our model as well.⁴

In the second part of the current investigation, we found that our simulation results were congruent with cross-task overall correlational analyses as well as item-by-item analyses of SD patient data reported in the literature (e.g., Blazely et al., 2005; Bozeat et al., 2000; Graham et al., 1994; Garrard & Carroll, 2006; Hodges et al., 1995; McKay et al., 2007). Specifically, the model's decline in lexical decision performance with semantic damage significantly correlated with its decline on all other tasks – both semantic and lexical; however, lexical decision exhibited an item-specific relationship only with word spelling and not with picture naming and other semantic tasks.

The reason why lexical decision performance declines with semantic damage but not necessarily on the same items as conceptual tasks is because the two are governed by two orthogonal dimensions, one captured by our measure of spelling consistency (consistency in mapping between orthographic and phonological features) and the other captured by our measure of concept consistency (consistency in mapping between visual and action/encyclopedic features). We investigated these two dimensions in our model and found that lexical decision performance co-varied with spelling consistency, whereas performance in

⁴It is worth pointing out that lexical decision is commonly tested with two distinct paradigms – the yes/no version, which was used in the imaging studies by Fujimaki et al. (2009) and Hauk et al. (2006), and the 2AFC version, which has been used with semantic dementia patients. The advantage of the 2AFC paradigm is that it allows the direct measure of the participants' preference between two possible letter strings. However, the yes/no paradigm has also been widely used and doesn't have some of the stimulus restrictions that the 2AFC version seems to require. Since our goal was to simulate the pattern of performance observed in semantic dementia, we implemented only the 2AFC paradigm. The implementation of the yes/no version may be somewhat different, and we haven't yet attempted it in our model. Nonetheless, the implantation of the 2AFC task is a substantial achievement, especially given the fact that the stimuli are homophonic pairs which differ in only a single letter – a task that would be difficult for any model, and indeed for most human participants.

all semantic tasks covaried with concept consistency; reading and spelling, on the other hand, co-varied with both. An interaction, whereby the advantage of semantically consistent items over inconsistent ones is especially pronounced when the items have inconsistent spelling, was observed in all tasks but lexical decision. Thus, we were able to explain seemingly contradictory findings in the literature – overall correlation between decline in lexical decision and in conceptual tasks, but no item -specific correlation (except perhaps for semantic tasks that involve written words). This is not because lexical decision deficits are due to damage to a lexicon – distinct from but neighboring the semantic system; rather it is due to damage in **a common integrative representation that subserves both semantic and lexical knowledge!** Together with previous neuropsychological and computational projects reconciling the findings of association and dissociation between semantic and lexical deficits in semantic dementia (Dilkina et al., 2008, Patterson et al., 2006, Plaut, 1997, Woolams et al., 2007), the results of the present research, – which address the challenge of lexical decision – lend strong support to the parsimonious notion that lexical and semantic processing depend on a single, common representation and processing system. Damage to this system, as seen in patients with semantic dementia, leads to a specific neuropsychological profile characterized by conceptual and lexical impairment with marked sensitivity to typicality and consistency.

4. Experimental procedure

Network architecture

The overall model architecture is the same as the connectionist model we previously used to account for naming and reading deficits in semantic dementia (Dilkina et al., 2008). There are four input/output (a.k.a. visible) layers, the two hidden layers, and a control layer consisting of six units that can be turned on or off to regulate which of the layers participate in a given task (Figure 1). This is achieved by a very strong negative bias on all units in the six processing layers so that unless the control units are activated, these layers remain insensitive to inputs and do not participate in processing or learning. Each of the six control units is fully connected to one of the six processing layers. Activating a control unit raises the resting level of the corresponding participating layer up to -3.00 , which then allows further excitatory input to bring this layer into play during processing. Both the inhibitory bias and the excitatory control connection weights are hard-coded.

Simulation materials

The current simulations used an extended version of the materials used by Dilkina et al. (2008), which consisted of 240 items from 12 categories. These materials and the algorithms used to obtain them are described in detail in the original paper (see Appendix A and B, Dilkina et al., 2008). Each item had a visual, an action/encyclopedic, an orthographic, and a phonological pattern. As in the previous work, the visual patterns are 60-item long binary vectors corresponding to a visual representation of the object. The original 60-item long action patterns corresponding to representations of how one interacts with the objects were extended to 120-item long action/encyclopedic patterns which also represent encyclopedic knowledge (e.g., where an animal lives, how an appliance works, etc.; cf. Rogers et al., 2004a). Essentially, these were two separate 60-item long binary vectors – one for action representations and one for encyclopedic representations. However, they were combined so that each item had a unique representation over every input/output layer (Note that action patterns alone were not necessarily unique). The 240 visual, action, and encyclopedic patterns were generated from probabilistic prototypes for each of the 12 item categories. As with the visual and action patterns borrowed from Dilkina et al. (2008), the encyclopedic prototypes were based on human ratings and similarity judgment data, and the individual patterns were created using a procedure similar to that used in Rogers et al. (2004a).

The phonological and orthographic representations were created in the same way as those used in Dilkina et al. (2008). They had a CVCC structure (C = consonant, V = vowel) designed to approximate English spelling and pronunciation co-occurrence statistics. In particular, we wanted the lexical items not only to exhibit the one-to-many mapping between graphemes and phonemes which is characteristic for the English language, but also to reflect the asymmetry in mapping letter s-to-sounds versus sound s-to-letters (the latter showing a stronger one-to-many pattern than the former). The ‘words’ were generated from 12 possible consonants with matching graphemes and phonemes and vowels which formed four groups of two phonemes and four graphemes (Table 2a), adding to a total of 16 possible vowel graphemes and 8 possible vowel phonemes. Thus, the only irregularities between spelling and pronunciation were in the vowels. Half of the vowel *graphemes* (occurring in about 15% of the words) had only one possible pronunciation, while the other half had two possible pronunciations. On the other hand, half of the vowel *phonemes* (occurring in about 55% of the words) had two possible spellings, while the other half had four possible spellings. The correspondence of graphemes to phonemes is called **regularity** (words can have regular or irregular pronunciation), whereas the correspondence of phonemes to graphemes is called **consistency** (words can have consistent or inconsistent spelling). Given a certain spelling, only one possible pronunciation is regular; the rest are considered irregular. Similarly, given a certain pronunciation, only one possible spelling is consistent; the rest are considered inconsistent.

The four vowel groups included six types of items, each of which was further divided into high frequency (HF) and low frequency (LF) items. The exact number of each type of item in a group is shown in Table 2b. These numbers were based on a corpus analysis of about 50,000 spoken word lemmas from the Celex English Lemma Database (Burnage, 1990), as described in Dilkina et al. (2008).

The lexical patterns were randomly matched with the visual and the action/encyclopedic patterns to produce 240 specific items each with four patterns – visual, action/encyclopedic, phonological, and orthographic.

Network training

Training consisted of a series of pattern presentations, which lasted for seven simulated unit time intervals each. During the first three intervals, an input pattern corresponding to the item being processed is clamped onto the appropriate layer. For example, if the trial requires mapping from visual input to action/encyclopedic output, then the visual pattern of the relevant item is clamped on. In addition, the control layer is also clamped to indicate what processing layers need to be used to accomplish the task. For example, in this same trial that requires mapping from visual input to action/encyclopedic output, the control units for the integrative layer and the action/encyclopedic layer will be on while all other control units will be off (including the one for the visual layer). For the remaining four intervals, the input is removed, and the network is allowed to adjust the activation of all units in all layers, including the one(s) previously clamped. During the final two intervals, the activations of units are compared to their corresponding targets.

The presence of control or task units is common in connectionist networks, especially when they are as large as the one presented here. The main reason for the use of such units is to assist the network in utilizing the available pathways in task-specific ways. If all layers, and therefore all sets of connections, are available to participate during the learning of the mapping between any pair of input/output patterns, then all connections are recruited for all mappings – because this is the fastest way to reduce the large amount of error present at the beginning of training. This is unadvisable in a large system, which has to learn the associations among a number of different surface forms. The control units are used in order

to encourage the network to selectively recruit subsets of the available pathways based on the task at hand.

The model was trained progressively through two stages designed to approximate developmental experience. During the first stage, the network was trained on visual, action/encyclopedic, and phonological – but not orthographic – knowledge. The integrative layer was used for all of these mappings, while the direct pathway was not used at all. During the second stage, the network learned to map among all four types of knowledge, including orthographic, through the integrative layer. Trials that involved only phonology and orthography used either the direct pathway alone (a third of the time) or the direct pathway in addition to the integrative layer (two thirds of the time).⁵

The network was trained on two types of mappings: one-to-one and one-to-all mappings. Both of these were included in order to allow the network to learn about all the possible associations among the four surface representations in a variety of contexts (under different control conditions). The one-to-one trials involve a single input pattern and a single target. An example of a one-to-one mapping is having the visual pattern for an item as input and the phonological pattern for that item as target. The one-to-all trials involve a single input and all four targets. In stage 1, one-to-all trials involve a single input and three targets (visual, action/encyclopedic, and phonological; remember that there is no orthographic training in this stage).

In each stage, the possible inputs were always seen in equal ratio, while this was not the case for the targets – the orthographic target was seen 1/3 as often as all the other targets (with the assumption that producing written output is generally much less common than producing spoken output; the 1:3 ratio was a mild approximation of that). Table 3 outlines the specific distribution of trials in the two stages.

A frequency manipulation was applied to all training stages so that high frequency items were seen five times more often than low frequency items. Also, even though the training stages were blocked, the different training trials within a stage were not. The network was trained on all items and mappings within a stage in an interleaved manner and the order of the items was random.

The connection weights were updated after every example using back-propagation with standard gradient descent and no momentum, with a learning rate of 0.001 and a weight decay of 0.0000016. Training through each stage continued until the error asymptoted.

Network testing

In addition to lexical decision, the network was tested on multiple semantic tasks used with SD patients as well as two other popular lexical tasks – word reading and word spelling. In each case, as during training, the task-specific input and controls were presented for the

⁵As explained by Dilkina et al. (2008), reading needed to be distributed over these two types of trials for technical reasons. At the beginning of stage 2, the connections between the integrative and the phonological layer have already grown in size due to the learning in stage 1. As a result, during subsequent trials involving phonological input or output, those connections most heavily contribute to the activation – even if the direct layer also participates. Therefore, the semantic connections mediating between orthography and phonology become responsible for the majority of the error and continue to learn, while the direct pathway connections remain small and insignificant. To ensure that both pathways contribute to the mapping between orthography and phonology in stage 2, the direct pathway needs to be trained by itself.

⁶The weight decay in the direct pathway – i.e. between orthography and the direct layer and between phonology and the direct layer, as well as the self-connections within the direct layer – was half of the weight decay in the rest of the network. The reason for this smaller value was purely technical. Because the proportion of trials which involve the direct pathway is relatively small, learning in that pathway is stalled when the weight decay is large (What happens is that the learned connection weights on trials which require the direct pathway decay away during trials which do not require it).

initial three intervals of processing, then the input was taken away and the network was allowed to continue settling for the remaining four intervals. Here we describe the materials used for each task and how performance was assessed.

Naming—Trials consisted of a single input presentation – the visual pattern of an item along with the relevant control pattern as seen during training (i.e., only the integrative and phonological control units were on). All 240 items were used. At the end of each trial, the response at the phonological layer was determined by selecting the most active units at each of the onset, vowel, and coda positions. If the most active unit at any one of the onset, vowel, and coda positions had an activation value below 0.3, this was considered a ‘no response’. Otherwise, the response was categorized as either correct or incorrect depending on whether the full CVCC was accurately produced.

Word-picture matching (WPM)—Trials consisted of 10 separate input presentations. Only the integrative and action/encyclopedic controls were on. The first input was a spoken word cue presented at the phonological layer. The other nine were the picture alternatives for multiple-choice selection – presented at the visual layer. One of the alternatives was the target, namely the visual pattern corresponding to the phonological pattern presented as cue. The remaining eight alternatives were within-category foils.

All 240 items were used as cue/target. They appeared as foils 8 ± 5.66 times on average (range: 0–34; median: 8; Table 4a). At the end of each stimulus presentation (cue, target, or foil), the activation over the integrative layer was recorded. At the end of the trial, the pattern of integrative activation to the cue was compared to the pattern of activation in response to each of the nine alternatives. The cosine (i.e. normalized dot product) of each pair of vectors was used as a measure of similarity. The alternative which achieved highest similarity with the cue was selected, and the response was categorized as correct or incorrect based on whether the chosen item matched the phonological cue. If there was no highest value (i.e. similarity values with the cue were equal among alternatives), performance was considered to be random.

Pyramids & Palm Trees (PPT)—This task was designed to approximate the PPT task done with patients (Howard & Patterson, 1992). Unfortunately, our model does not have explicit associative knowledge between items. The closest approximation is the encyclopedic representation, and so we designed our task trials based on item similarity over these representations.

Trials consisted of three separate input presentations – cue, target and foil. As with WPM, only the integrative and action/encyclopedic controls were on. All three stimuli were presented either at the visual layer (pictorial version of the task) or at the orthographic layer (word version of the task). All 240 items were used as cue. They appeared as target or foil 2 ± 1.31 times on average (range: 0–6; median: 2). As in the real PPT task, the two alternatives for multiple-choice selection were from a shared category, which was different from the category of the cue. They were selected to have an approximately equal *visual* similarity – but different *action/encyclopedic* similarity – with the cue (Table 4b). The item with higher action/encyclopedic similarity was the target, while other one was the foil. Performance was assessed as in WPM: the alternative with higher integrative pattern similarity with the cue was selected, and the response was categorized as correct (target), incorrect (foil), or random (when the target and the foil produced identical similarity with the cue).

Word reading—Trials consisted of a single input presentation – the orthographic pattern of an item along with the relevant control pattern as seen during training (i.e., the

integrative, direct, and phonological control units were on). All 240 items were used (36 HF regular, 180 LF regular, 8 HF irregular, 16 LF irregular). As for naming, at the end of each trial, the response at the phonological layer was determined by selecting the most active units at each of the onset, vowel, and coda positions. The response was categorized as either correct or incorrect depending on whether the full CVCC was accurately produced.

Word spelling—Trials consisted of a single input presentation – the phonological pattern of an item along with the relevant control pattern as seen during training (i.e., the integrative, direct, and orthographic control units were on). All 240 items were used (28 HF consistent, 156 LF consistent, 16 HF inconsistent, 40 LF inconsistent). At the end of each trial, the response at the orthographic layer was determined by selecting the most active units at each of the onset, vowel, and coda positions. The response was categorized as either correct or incorrect depending on whether the full CVCC was accurately produced.

Two-choice lexical decision (LD)—This task was designed to approximate the versions administered to the patients reported by both Patterson and colleagues (2006) and Blazely and colleagues (2005). Trials consisted of two separate input presentations – word and pseudohomophone. Both were presented at the orthographic layer. Performance was assessed under three different control conditions: (1) with the integrative and orthographic controls turned on; (2) with the direct and orthographic controls turned on; and (3) with the integrative, direct and orthographic controls turned on. The orthographic form of the nonword in each trial differed from the word it was paired with *only in the vowel*, and the alteration was such that the phonological pattern of the word and the nonword were identical, i.e. they were homophones. Thus, if the word was spelled CVCC and pronounced cvcc, then the nonword was spelled CV₁CC and pronounced cvcc. The pronunciation of the nonword was determined following the regular mappings between vowel graphemes and phonemes (see Table 2). Thus, it was not possible to create a matching pseudohomophone for all words as some regular words did not have another possible regular spelling of their vowel; we could only create pseudohomophones for words that were either both irregular and inconsistent (the pseudohomophone being regular and consistent) or both regular and consistent (the pseudohomophone being regular and inconsistent). There was a total of 116 trials (60 pairs where the spelling of the word was consistent, while the spelling of the nonword was inconsistent, W>NW; and 56 pairs where the spelling consistency was reversed, W<NW).

LD performance was assessed using a measure called *orthographic echo*. At the end of each stimulus presentation, the activation over the orthographic layer was recorded (Note that this is a *re*-activation of the pattern presented as input, hence the term *echo*). At the end of the trial, the strength of the orthographic echo of word and nonword were calculated by summing the activation values of the units corresponding to the four graphemes for the pair of items (CVCC for the word, CV₁CC for the nonword). The item with stronger orthographic echo was selected, and the response was categorized as correct (word) or incorrect (nonword).⁷

⁷We chose to use orthographic echo as our measure of lexical decision (instead of other measures used in the literature, e.g. stress or polarity over the semantic/integrative layer) for a number of reasons. First of all, this measure was very well suited for this 2AFC version of the task with homophonic word-nonword pairs, where performance requires making fine distinctions at the orthographic level. Secondly, we wanted to investigate the model's lexical decision performance in a variety of settings (with contribution from semantics and without); thus, implementing our measure at the orthographic layer allowed us to have a consistent assessment of performance across the different conditions. Finally, orthographic echo is a measure of the strength of re-activation of a specific written form; it looks at the word features presented at input, instead of looking at the entire layer. That way, it gives an accurate and precise measure of how interaction between the orthographic layer and the rest of the system may weaken the representation of the original stimulus (rather than looking at how this interaction may also strengthen other word representations).

In summary, our measure for lexical decision was orthographic echo, and we compared performance across three different contexts – when only the integrative layer was allowed to participate (i.e. to maintain activation over the orthographic layer), when only the direct layer was allowed to participate, and when both hidden layers were allowed to participate.

To simulate semantic dementia and its progressive nature, the network was damaged by removing integrative units and increasing the weight decay of integrative connections in combination with retraining. During training, the weight decay was 10^{-6} . During damage, 10% of the integrative units were removed and the weight decay of the connections among the remaining integrative units was doubled; the network was further trained with one pass through the entire training set, after which another 10% of the integrative units were removed and the weight decay was again doubled, then the network was further trained; etc. This was done 10 times – until there were no more units left in the integrative layer. Because after the last lesion there was no semantics left, performance on multiple choice tasks that require the integrative layer was guaranteed to be random. Thus, the 10 lesion levels used for the analysis are lesions 0 through 9, where lesion 0 is the performance of the network before damage.

To ensure appropriate sampling, four models were trained (identical architecture and training sets; only different random number generator seeds), and each was then damaged and tested five times (again, using different seeds) adding to a total of 20 different sets of results.

Research Highlights

- lexical decision correlated with picture naming, WPM, and PPT but not in item-specific ways
- LD also correlated with word reading and word spelling; only the latter was item-specific
- spelling consistency governs LD performance; concept consistency governs semantic performance; both factors affect reading and spelling
- semantics and the direct pathway contributed to LD performance in distinct but interactive ways

Acknowledgments

The work reported in this paper was partially supported by NIH Grant P01-MH64445.

References

- Azuma T, Van Orden GC. Why SAFE is better than FAST: The relatedness of a word's meaning affects lexical decision times. *Journal of Memory & Language*. 1997; 36(4):484–504.
- Benedet M, Patterson K, Gomez-Pastor I, de la Rocha MLG. 'Non-semantic' aspects of language in semantic dementia: As normal as they're said to be? *Neurocase*. 2006; 12:15–26. [PubMed: 16517512]
- Binder JR, McKiernan KA, Parsons ME, Westbury CF, Possing ET, Kaufman JN, Buchanan L. Neural correlates of lexical access during visual word recognition. *Journal of Cognitive Neuroscience*. 2003; 15:372–393. [PubMed: 12729490]
- Blazely A, Coltheart M, Casey BJ. Semantic impairment with and without surface dyslexia: Implications for models of reading. *Cognitive Neuropsychology*. 2005; 22:695–717. [PubMed: 21038273]

- Bozeat S, Lambon Ralph MA, Patterson K, Garrard P, Hodges JR. Non-verbal semantic impairment in semantic dementia. *Neuropsychologia*. 2000; 38:1207–1215. [PubMed: 10865096]
- Burnage, G. CELEX English lexical user guide. In: Burnage, G., editor. CELEX – A Guide for Users. Nijmegen: Centre for Lexical Information, University of Nijmegen; 1990.
- Caramazza A. How many levels of processing are there? *Cognitive Neuropsychology*. 1997; 14:177–208.
- Coltheart M, Rastle K, Perry C, Langdon RJ, Ziegler JC. DRC: A dual route cascaded model of visual word recognition and reading aloud. *Psychological Review*. 2001; 108:204–256. [PubMed: 11212628]
- Dell GS, O'Seaghdha PG. Stages of lexical access in language production. *Cognition*. 1992; 42:287–314. [PubMed: 1582160]
- Dilkina K, McClelland JL, Plaut DC. A single-system account of semantic and lexical deficits in five semantic dementia patients. *Cognitive Neuropsychology*. 2008; 25:136–164. [PubMed: 18568816]
- Forster KI, Hector J. Cascaded versus noncascaded models of lexical and semantic processing: The turtle effect. *Memory & Cognition*. 2002; 30:1106–1117.
- Funnell E. Response biases in oral reading: An account of the co-occurrence of surface dyslexia and semantic dementia. *The Quarterly Journal of Experimental Psychology*. 1996; 49A:417–446. [PubMed: 8685392]
- Fujimaki N, Hayakawa T, Ihara A, Wei Q, Munetsuna S, Terazono Y, Matani A, Murata T. Early neural activation for lexico-semantic access in the left anterior temporal area analyzed by an fMRI-assisted MEG multidipole method. *Neuroimage*. 2009; 44:1093–1102. [PubMed: 19027078]
- Garrard P, Carroll E. Lost in semantic space: A multi-modal, non-verbal assessment of feature knowledge in semantic dementia. *Brain*. 2006; 129:1152–1163. [PubMed: 16585052]
- Graham KS, Hodges JR, Patterson K. The relationship between comprehension and oral reading in progressive fluent aphasia. *Neuropsychologia*. 1994; 32:299–316. [PubMed: 8202225]
- Graham NL, Patterson K, Hodges JR. The impact of semantic memory impairment on spelling: Evidence from semantic dementia. *Neuropsychologia*. 2000; 38:143–163. [PubMed: 10660226]
- Hauk O, Patterson K, Woollams A, Watling L, Pulvermuller F, Rogers TT. [Q:] When would you prefer a SOSSAGE to a SAUSAGE? [A:] At about 100 msec. ERP correlates of orthographic typicality and lexicality in written word recognition. *Journal of Cognitive Neuroscience*. 2006; 18:818–832. [PubMed: 16768380]
- Hodges JR, Graham N, Patterson K. Charting the progression of semantic dementia: Implications for the organisation of semantic memory. *Memory*. 1995; 3:463–495. [PubMed: 8574874]
- Howard, D.; Patterson, K. *Pyramids and Palm Trees: A Test of Semantic Access from Pictures and Words*. Bury St. Edmunds, UK; Thames Valley Publishing Company: 1992.
- Jefferies E, Lambon Ralph MA, Jones R, Bateman D, Patterson K. Surface dyslexia in semantic dementia: A comparison of the influence of consistency and regularity. *Neurocase*. 2004; 10:290–299. [PubMed: 15788266]
- Knibb JA, Kipps CM, Hodges JR. Frontotemporal dementia. *Current Opinions in Neurology*. 2006; 19:565–571.
- Knott R, Patterson K, Hodges JR. Lexical and semantic binding effects in short-term memory: Evidence from semantic dementia. *Cognitive Neuropsychology*. 1997; 14:1165–1216.
- Knott R, Patterson K, Hodges JR. The role of speech production in auditoryverbal short-term memory: Evidence from progressive fluent aphasia. *Neuropsychologia*. 2000; 38:125–142. [PubMed: 10660225]
- Kuchinke L, Jacobs AM, Grubich C, Võ MLH, Conrad M, Herrmann M. Incidental effects of emotional valence in single word processing: An fMRI study. *NeuroImage*. 2005; 28:1022–1032. [PubMed: 16084739]
- Lambon Ralph MA, Graham KS, Patterson K. Is picture worth a thousand words? Evidence from concept definitions by patients with semantic dementia. *Brain & Language*. 1999; 70:309–335. [PubMed: 10600223]
- Lambon Ralph MA, Howard D. Gogi aphasia or semantic dementia? Simulating and assessing poor verbal comprehension in a case of progressive fluent aphasia. *Cognitive Neuropsychology*. 2000; 17:437–465. [PubMed: 20945190]

- Levelt, WJM. Speaking, from Intention to Articulation. Cambridge: MIT Press; 1989.
- Marslen-Wilson WD. Functional parallelism in spoken word recognition. *Cognition*. 1987; 25:71–102. [PubMed: 3581730]
- McKay A, Castles A, Davis C, Savage G. The impact of progressive semantic loss on reading aloud. *Cognitive Neuropsychology*. 2007; 24:162–186. [PubMed: 18416487]
- Moss HE, Tyler LK. Investigating semantic memory impairments: The contribution of semantic priming. *Memory*. 1995; 3:359–395. [PubMed: 8574870]
- Neary D, Snowden JS, Gustafson L, Passant U, Stuss D, Black S, et al. Frontotemporal lobar degeneration. *Neurology*. 1998; 51:1546–1554. [PubMed: 9855500]
- Papagno C, Capitani E. Slowly progressive aphasia: A four-year follow-up study. *Neuropsychologia*. 2001; 39:678–686. [PubMed: 11311298]
- Patterson K. The reign of typicality in semantic memory. *Philosophical Transactions of the Royal Society B: Biological Sciences*. 2007; 362:813–821.
- Patterson K, Lambon Ralph MA, Jefferies E, Woollams A, Jones R, Hodges JR, Rogers TT. “Presemantic” cognition in semantic dementia: Six deficits in search of an explanation. *Journal of Cognitive Neuroscience*. 2006; 18:169–183. [PubMed: 16494679]
- Pexman PM, Hargreaves IS, Edwards JD, Henry LC, Goodyear BG. The neural consequences of semantic richness. *Psychological Science*. 2007; 18:401–406. [PubMed: 17576279]
- Plaut DC. Double dissociation without modularity: Evidence from connectionist neuropsychology. *Journal of Clinical and Experimental Neuropsychology*. 1995; 17:291–321. [PubMed: 7629273]
- Plaut DC. Structure and function in the lexical system: Insights from a distributed model of word reading and lexical decision. *Language & Cognitive Processes*. 1997; 12:765–805.
- Plaut DC, McClelland JL, Seidenberg M, Patterson K. Understanding normal and impaired word reading: Computational principles in quasi-regular domains. *Psychological Review*. 1996; 103:56–115. [PubMed: 8650300]
- Rochon E, Kavé G, Cupit J, Jokel R, Winocur G. Sentence comprehension in semantic dementia: A longitudinal case study. *Cognitive Neuropsychology*. 2004; 21:317–330. [PubMed: 21038208]
- Rosch E, Mervis CB, Gray WD, Johnson DM, Boyes-Braem P. Basic objects in natural categories. *Cognitive Psychology*. 1976; 8:382–439.
- Rogers TT, Lambon Ralph MA, Garrard P, Bozeat S, McClelland JL, Hodges JR, et al. The structure and deterioration of semantic memory: A neuropsychological and computational investigation. *Psychological Review*. 2004a; 111:205–235. [PubMed: 14756594]
- Rogers TT, Lambon Ralph MA, Hodges JR, Patterson K. Natural selection: The impact of semantic impairment on lexical and object decision. *Cognitive Neuropsychology*. 2004b; 21:331–352. [PubMed: 21038209]
- Saffran EM, Coslett HB, Martin N, Boronat CB. Access to knowledge from pictures but not words in a patient with progressive fluent aphasia. *Language & Cognitive Processes*. 2003; 18:725–757.
- Samson D, Pillon A. Orthographic neighborhood and concreteness effects in the lexical decision task. *Brain & Language*. 2004; 91:252–264. [PubMed: 15485714]
- Schwartz MF, Marin OSM, Saffran EM. Dissociation of language function in dementia: A case study. *Brain & Language*. 1979; 7:277–306. [PubMed: 455049]
- Tyler LK, Moss HE. Going, going, gone...? Implicit and explicit tests of conceptual knowledge in a longitudinal study of semantic dementia. *Neuropsychologia*. 1998; 36:1313–1323. [PubMed: 9863685]
- Tyler LK, Moss HE, Galpin A, Voice JK. Activating meaning in time: The role of imageability and form-class. *Language & Cognitive Processes*. 2002; 17:471–502.
- Woollams AM, Lambon Ralph MA, Plaut DC, Patterson K. SD-squared: On the association between semantic dementia and surface dyslexia. *Psychological Review*. 2007; 114:316–339. [PubMed: 17500629]
- Wurm LH, Vakoch DA, Seaman SR. Recognition of spoken words: Semantic effects in lexical access. *Language & Speech*. 2004; 47:175–204. [PubMed: 15581191]

Figure 1.

Architecture of the connectionist model of semantic and lexical processing: dashed lines signify connections with large hard-coded positive weights from the units in the control layer to each of the six processing layers; solid lines signify connections with small randomly-initiated weights, which are learned over the course of training; arrows indicate direction of connectivity.

Figure 2. Lexical decision performance on W>NW pairs and W<NW pairs as a function of lesion extent.

Figure 3. Effect of spelling typicality in the two pathways involved in lexical decision. NOTE: log odds of 0 marks chance level.

Figure 4. Correlations between lexical decision and other semantic and lexical tasks performed by the network.

Figure 5. Effects of spelling consistency and concept consistency on the seven tasks in the model. Effect = performance on consistent items – performance on inconsistent items. (5a) main effect of spelling consistency; (5b) main effect of concept consistency; (5c) interaction between spelling consistency and concept consistency. (***) $p < .0005$; (**) $p < .005$; (*) $p < .05$)

Table 1

Item-specific correlations among tasks.

	WPM	PPTp	PPTw	reading	spelling	lexical decision	W<NW	lexical decision
naming	***	***	***	***	***			
WPM		**	***	***	***			
PPTp			***		*			
PPTw				*	*			
reading					***			
spelling						**		***

 $p < .0005$
**
 $p < .005$
*
 $p < .05$

Table 2

(a) Vowel phoneme-grapheme combinations.					
	grapheme 1	grapheme 1	grapheme 3	grapheme 4	
phoneme 1	regular, consistent	irregular, inconsistent			
phoneme 2	irregular, inconsistent	regular, consistent	regular, inconsistent	regular, inconsistent	

(b) Number of HF and LF items in each of the four groups of vowel phoneme-grapheme combinations.							
	grapheme 1		grapheme 2		grapheme 3		grapheme 4
	HF	LF	HF	LF	HF	LF	LF
phoneme 1	4	27	1	1			
phoneme 2	1	3	3	12	1	5	1

Table 3

Trial types and associated controls (i.e., active layers) used in the two stages of network training.

stage	input	target(s)	controls	proportion of total trials
stage 1	V	V	VI	1/12
		E	EI	1/12
		P	PI	1/12
		VEP	VEPI	1/12
	E	V	VI	1/12
		E	EI	1/12
		P	PI	1/12
		VEP	VEPI	1/12
	P	V	VI	1/12
		E	EI	1/12
		P	PI	1/12
		VEP	VEPI	1/12
stage 2	V	V	VI	9/156
		E	EI	9/156
		O	OI	3/156
		P	PI	9/156
		VEOP	VEOPI	9/156
	E	V	VI	9/156
		E	EI	9/156
		O	OI	3/156
		P	PI	9/156
		VEOP	VEOPI	9/156
	O	V	VI	9/156
		E	EI	9/156
		O	OID	2/156
			OD	1/156
		P	PID	6/156
			PD	3/156
		VEOP	VEOPID	9/156
	P	V	VI	9/156
		E	EI	9/156
		O	OID	2/156
			OD	1/156
		P	PID	6/156
			PD	3/156
		VEOP	VEOPID	9/156

abbreviations: V = visual; E = action/encyclopedic; O = orthographic; P = phonological; I = integrative; D = direct.

Table 4

(a) Similarity of the cue to the foils in WPM, measured as the cosine of the two patterns.				
	visual similarity (cue, foil)	action/encyclopedia similarity (cue, foil)	visual and action /encyclopedia similarity (cue, foil)	
average (± st dev)	.66 ± .10	.62 ± .11	.68 ± .09	
range	.29 to .88	.30 to .82	.46 to .86	

(b) Similarity of the cue to the target and foil in PPT, measured as the cosine of the two patterns.				
	visual similarity (cue, target)	visual similarity (cue, foil)	difference in visual similarity	action/encyclopedia similarity (cue, target)
average (± st dev)	.33 ± .08	.32 ± .08	.01 ± .03	.42 ± .14
range	.21 to .53	.20 to .51	-.06 to .10	.17 to .78
				action/encyclopedia similarity (cue, foil)
				.17 ± .13
				.00 to .55
				difference in action/encyclopedia similarity
				.25 ± .06
				.14 to .46