Coherent concepts are computed in the anterior temporal lobes

Matthew A. Lambon Ralph^{a,1}, Karen Sage^a, Roy W. Jones^b, and Emily J. Mayberry^a

^aNeuroscience and Aphasia Research Unit, School of Psychological Sciences, University of Manchester, Manchester M13 9PL, United Kingdom; and ^bResearch Institute for the Care of Older People, Bath BA1 3NG, United Kingdom

Edited by Edward E. Smith, Columbia University, New York, NY, and approved December 24, 2009 (received for review July 1, 2009)

In his *Philosophical Investigations*, Wittgenstein famously noted that the formation of semantic representations requires more than a simple combination of verbal and nonverbal features to generate conceptually based similarities and differences. Classical and contemporary neuroscience has tended to focus upon how different neocortical regions contribute to conceptualization through the summation of modality-specific information. The additional yet critical step of computing coherent concepts has received little attention. Some computational models of semantic memory are able to generate such concepts by the addition of modality-invariant information coded in a multidimensional semantic space. By studying patients with semantic dementia, we demonstrate that this aspect of semantic memory becomes compromised following atrophy of the anterior temporal lobes and, as a result, the patients become increasingly influenced by superficial rather than conceptual similarities.

conceptual knowledge | semantic dementia | semantic memory

In neuroscience, semantic memory or conceptual knowledge refers to the font of information we possess that enables us to bring meaning to words, objects, and all other nonverbal stimuli (e.g., smell, sounds, etc.). When this knowledge breaks down or becomes inaccessible, as it does in a range of neurological conditions, patients are left with significant disability. Key questions for neuroscience, therefore, include: What brain regions contribute to semantic memory and what role do they play? A potential answer can be found in both the classical and contemporary neuroscience literature. Wernicke and Meynert (see ref. 1) proposed that conceptualization reflected the conjoint action of the different cortical association areas that code sensory, motor, and verbal information (engrams). Thus, via a web of interconnections (see Fig. 1A), it is possible to taste an apple and in turn automatically activate all of the other knowledge associated with it (e.g., what it looks like, its internal and external color, how it is prepared in various culinary dishes, what it is called, etc.). A very similar idea is still the dominant one in contemporary neuroscience. Modern neuroimaging has been able to show that multiple association regions are activated when participants are required to process the meaning of different concepts (2, 3). It is also appealing because there is no mystery about where semantic knowledge comes from—it derives directly from the totality of our verbal and nonverbal experience. In turn, unlike classical symbolic approaches, this knowledge can be retrieved by reactivating the representation or engram of this experiential database (4).

The issue at the center of this study is how this multimodal experience can be brought together to form “coherent” and generalizable semantic representations (5, 6). Whereas the source of this information is obvious, the computations required are non-trivial because these semantic features or atoms combine in complex, nonlinear ways. Some of the computational challenges include the following (illustrated with respect to the objects shown in Fig. 2):

- (i). *Convergence of modality-specific and event-specific information*: Potentially any sensory, motor, and verbal modalities can contribute to our semantic knowledge, and this information does not necessarily arise at the same point in

time (7). Thus, a mechanism is needed that systematically draws all this information together and does so in a time-invariant fashion. For example, the everyday objects in Fig. 2 can be coded in terms of praxic, visual, auditory, and verbal attributes but the important similarities and differences can only be computed when this information is combined, even though this experience is extended over multiple episodes and key information about an item may never co-occur in the same episode (e.g., we never clean and boil kettles at the same time).

- (ii). *Features span different ranges of concepts*: Both verbal and nonverbal features can extend in different, unrelated ways. For example, all of the items in Fig. 2 are designed to hold and pour liquids and can be called “containers.” Some are designed solely for wine or oil (but not necessarily the same type of oil or wine), and even among the water containers there are important variations in terms of the temperature and function of each one.
- (iii). *Complex set of nonlinear relationships*: As an extension of (ii), it is clear that the multitude of features does not align and can sometimes be orthogonally or nonlinearly related. For example, the construction material (e.g., copper) is not related to the object’s core function, although it does relate to the way in which it should be handled (copper needs to be cleaned, china is fragile, etc.). Additionally, a core aspect of being a kettle is that water (or sometimes other liquids) are heated in them but this feature varies for teapots; for some (those made of plastic or china), direct heating would lead to undesirable consequences, but for others it is necessary (where infusion takes a long time).
- (iv). *Surface similarities are only a partial guide to meaning*: Surface (in this case visual) similarities are not a perfect guide to object identity. In fact, a range of opposing problems are captured by Fig. 2: (a) some prominent aspects of an object are superfluous—for example, possessing a floral design; (b) some aspects (handles and spouts) are common across the broad range of concepts but, as a consequence, finer categorization of function (e.g., teapots vs. kettles) requires sensitivity to subtle visual variations; and (c) however, this same category membership also has to be extended to exemplars that are visually very different (the animal-shaped objects).
- (v). *Semantic-based generalization to new or changing concepts*: Another critical ability is semantic generalization—we often encounter new exemplars of an object (e.g., a new

Author contributions: M.A.L.R., K.S., and E.J.M. designed research; R.W.J. and E.J.M. performed research; M.A.L.R. and E.J.M. analyzed data; and M.A.L.R., K.S., R.W.J., and E.J.M. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

¹To whom correspondence should be addressed. E-mail: matt.lambon-ralph@manchester.ac.uk.

This article contains supporting information online at www.pnas.org/cgi/content/full/0907307107/DCSupplemental.

Fig. 1. Diagrams of the distributed semantic hypothesis, the hub-and-spoke model, and examples of the ATL atrophy present in semantic dementia. In many traditional and contemporary accounts of conceptualization, semantic representations are formed as a by-product of the interaction between sensory, verbal, and motor association cortices. This notion is summarized in A. A revision of this idea is presented in B in which, in addition to modality-specific sources, conceptualization also involves modality-invariant representations. Through their interaction, this framework is able to address various key computational challenges that are posed when concepts are formed from transmodal experience (see text). In this study, we tested the hypothesis that this invariant contribution can break down after damage to the ATL by assessing a series of patients with semantic dementia, which produces a selective semantic impairment in the context of relatively localized atrophy of the ATL. Sample MR images from one patient (MT) are shown in C.

teapot) or existing examples changing over time (e.g., your favorite teapot gets chipped, the decoration fades, or you lose its lid), yet we easily and automatically generalize the knowledge about teapots to new or changing examples, even though we have never experienced these specific exemplars before.

Although primarily focused on language, these and many other challenges were highlighted by Wittgenstein in his *Philosophical Investigations* (8). Wittgenstein noted that experience and use play a key role in concepts, but that extracting shared meaning is not always possible on the basis of identifying a set of

Fig. 2. Example of the complex relationship between features and concepts. As can be seen, although some visual similarities and features are related (e.g., made of copper, requires specific cleaning), many are not (e.g., fragile). Instead, there is a complex tapestry of features, each of which extend over a different area. These complex arrangements can be mapped by projecting the modality-specific information into a high-dimensional, modality-invariant space (in the anterior temporal lobes). See text.

shared features. As noted above, this includes the observation that exemplars can vary in form (e.g., the different types of handle—break, switch, pump, crank—in a train driver’s cabin; §12) or that knowledge can be generalized to new exemplars (e.g., knowledge of how to use the king in a set of chess pieces even if the design changes; §31), and, most famously, that conceptually related items do not necessarily share any particular feature in common and thus cannot be defined in that way (e.g., games; §66). In considering different types of game in terms of various constituent features, Wittgenstein’s summary could equally apply to the objects in Fig. 2: “And we can go through the many, many other groups of games in the same way; can see how similarities crop up and disappear. And the result of this examination is: we see a complicated network of similarities overlapping and criss-crossing: sometimes overall similarities, sometimes similarities of detail.”

A potential solution comes in three parts: philosophical-cognitive science, computational, and neuroanatomical. Rather than searching endlessly for defining attributes for each and every set of concepts, some philosophers [e.g., Wittgenstein (§67) (8)] and cognitive scientists have proposed that “coherent representations” might follow if an additional computation or representation is added to our sensory-verbal experience (a full discussion of the various suggested solutions goes beyond the scope of this paper, but broad and clear reviews can be found in refs. 5, 6). This idea parallels a similar problem and approach found in computational modeling. Following Wittgenstein’s example, consider four category exemplars with the features {A,B,C}, {A,B,D} and {A,C,D}, {B,C,D}. No single feature defines them as a single group but instead they form a group through partial, overlapping features. In computational terms, this problem cannot be solved by a single layer of feature-coding units (a single-layer perceptron) because the representations are not linearly separable. They can, however, be grouped properly if an additional (hidden) layer is added because this allows rerepresentation of the feature input. This type of solution can be found in a number of implemented computational models of semantic memory, and is summarized graphically in Fig. 1B (7, 9, 10). In this “hub-and-spoke” framework, semantic representations are formed and activated through the joint action of different types of information. Like the classical and contemporary approaches to semantic memory (2–4), the spokes (each oval) represent the information that arises in each of our sensory, motor, and verbal association cortices. These are entry and exit points for both comprehension and expression of semantic knowledge whether in

the verbal or nonverbal domain. If these were only directly connected to each other, as per the Wernicke-Meynert model (1) (see Fig. 1A), then they would form a single-layer perceptron. As argued above, whereas this semantic perceptron might be able to generate some concepts through the amalgamation of features, it would not be able to deal with the computational challenges listed above.

These challenges can be met by the addition of an extra, intermediating representational layer (7, 9). This allows the formation of modality-invariant multidimensional representations that, through the cross-translation of information between each modality, code the higher-order statistical structure that is present in our transmodal experience of each entity. As such, the same core information is activated each time an entity is encountered even if different aspects occur in separate episodes (challenge 1); verbal and nonverbal features can be linked to different ranges of concept (challenge 2) whether or not the information is systematically or orthogonally related (challenge 3). As well as being able to code the partial similarity structure in specific domains, the modality-invariant representations add greater flexibility to deal with concepts that do not follow these surface similarities (challenge 4). And the system provides a mechanism for generalization to new or changing exemplars (challenge 5).

The purpose of this study was to test that this higher-order aspect of semantic memory can break down after brain damage and, moreover, that it is associated with the anterior temporal lobes (ATL). The ATL are a potentially suitable substrate for the neuroanatomical implementation of a modality-invariant hub because of two key features: (i) It is not associated with any single motor, sensory, or verbal input/output (which would otherwise dominate its function) (10) but (ii) it is widely connected to other temporal, parietal, and frontal regions (11, 12), allowing different streams of modality-specific information to converge on this point (13, 14). Although this region does not appear in classical models of aphasia (13–15), there is a growing consensus that this region is critically involved in semantic memory (15–17). This is based on findings that verbal and nonverbal semantic performance is impaired after ATL damage (16, 18, 19), is selectively slowed when repetitive transcranial magnetic stimulation is applied to the lateral ATL in normal participants (20, 21), and is activated in neuroimaging studies of semantic processing (22, 23). Although these studies clearly implicate the ATL in semantic memory, many questions about their exact contribution and role remain (15, 17). In this study, therefore, we tested the hypothesis that the ATL underpin the modality-invariant contribution to coherent representations.

The prediction, from the framework set out above, is that when the ATL modality-invariant information is damaged, not only will semantic processing be significantly compromised (because the hub is a major conduit for information flow between modality-specific centers) but performance will become dominated by the modality-specific surface similarities and be less reflective of higher-order semantic structure. We tested this prediction in a case series of patients with semantic dementia (the temporal lobe variant of frontotemporal dementia; see *Methods and Materials* for patient background data) (24). This progressive neurological disorder has three aspects which provide a unique opportunity to explore this aspect of semantic memory: (i) the disease produces relatively focused atrophy and associated hypometabolism of the ATL bilaterally (see Fig. 1C) (25, 26); (ii) the patients present with a strikingly selective impairment of semantic memory without concomitant impairments in other aspects of higher cognitive function which can interfere with task performance (15, 27); and (iii) although selective in nature, the disorder is progressive, giving the opportunity to investigate semantic performance at different levels of disease severity.

Results

To test the hypothesis that the ATL-semantic hub is critically important for coding higher-order modality-invariant structure, we created a semantic matching-to-sample task and tested this in six semantic dementia (SD) patients covering the range of mild to severe semantic impairment (see *Methods and Materials* for details on task design). To investigate the integrity of this higher-order structure, we selected samples (exemplars) so as to pit visual similarity against category membership. Everyday exemplars can vary a great deal in terms of their superficial similarities (e.g., handles in a train cabin, types of bird, types of teapot; Fig. 2). With careful selection it is possible to find everyday objects or animals that either do or do not conform to the surface similarity (Fig. 3). Not only is it possible to find both typical and less typical examples of the target (e.g., cat: burmese vs. sphynx) but also nontarget items that vary in their superficial similarity to the target (e.g., chihuahua vs. mink vs. train). The modality-invariant contribution to semantic memory should be especially important not only when extending category membership to an atypical exemplar (e.g., the hairless sphynx cat) but also when rejecting a “pseudotypical” noncategory member (e.g., the chihuahua). We expected to observe, therefore, two kinds of error in the SD patients’ performance: overgeneralizations (incorrectly selecting the pseudotypical noncategory member) and undergeneralizations (failing to select the atypical category exemplar). Furthermore, because *both* errors arise as a result of deterioration to the *same* category boundary (*Discussion*), on at least some occasions we expected the SD patients to make over- and under-generalization errors simultaneously.

Fig. 3. Semantic dementia performance on the matching-to-sample task. Performance for six SD patients and control participants is shown in the top panel. Error bars denote the SE of the control mean per condition. Asterisks denote abnormal performance. Examples of the five types of item within the choice array are shown below the relevant conditions (see *SI Text* for further details). Each trial contained a mixture of targets and foils. The targets included both typical and atypical examples of each target concept (e.g., cat, ball, fish, shoe). There were also three types of foil. Some were pseudotypical examples which, although not an exemplar of the target concept, share many features in common; some were partially related; and the remainder were unrelated. The SD patients made a combination of two types of selection error: (i) undergeneralization—failures to pick examples of the target concept (especially of the atypical exemplars) and (ii) overgeneralization—incorrect selection of nonconcept items (particularly pseudotypical and some partially related choices).

As can be seen in Fig. 3, the SD patients performed exactly as expected. Whereas performance dropped a little for typical targets and partially related foils, there was a much greater decline in performance on the two key conditions: atypical targets and pseudotypical foils [leading to a significant interaction between participant group (patients vs. controls) and choice type: $F(3,33) = 8.3$, $P < 0.001$]. The patients were ordered by target selection accuracy. Rejecting the pseudotypical foils proved to be a significant challenge for all patients; even the best-performing patient fell below the control range in this condition. With increasing severity, the patients began to overgeneralize the target concept to the partially related foils as well. The patients also undergeneralized the same concepts, particularly for the atypical targets. Although one patient performed within the normal range on this condition, performance on this condition dropped away steeply in line with semantic severity. Perhaps most strikingly, all of the patients made at least some simultaneous over- and undergeneralization of the same concept. On average, the patients achieved perfect selection on 41% of the trials. They simultaneously over- and undergeneralized on 18% of trials, they undergeneralized alone on 19% of trials; and overgeneralized alone on 22% of trials. In contrast, none of the patients had any difficulty in rejecting all of the unrelated choices—indicating that even the patients with severe semantic impairment were not selecting choices randomly. Finally, it is intriguing to note that overgeneralization seems to foreshadow undergeneralization in these data, perhaps suggesting that, at least in the early phase of semantic decline, conceptual boundaries inflate before contracting (like a dying star). Such direct comparisons need to be treated with caution, however, because it is difficult (if not impossible) to balance the intrinsic semantic difficulty in each condition—indeed, it would appear that on average the normal participants also tended to find the pseudotypical foils slightly harder to reject.

Discussion

We have a wealth of accumulated sensory, verbal, and motor experience of words, objects, and people. We can use it automatically to bring meaning to words and a variety of nonverbal stimuli (sounds, pictures, objects, etc.). We can, in turn, convey this knowledge through both language and nonverbal expression (e.g., drawing, using objects, etc.). When this knowledge—*semantic memory*—becomes impaired, as it does in a range of neurological conditions, patients are left with significant disability. Both classical and contemporary models of semantic memory (1, 2, 4) have advanced the intuitively appealing notion that semantic representations are constructed from the joint action of multiple, modality-specific association cortical regions, each of which codes the structure of information arising in a specific sensory, motor, or verbal domain. In this study, we tested and confirmed the hypothesis that, in addition to these modality-specific information sources, the formation of coherent concepts requires a set of modality-invariant representations and that these are underpinned by the anterior temporal lobes.

This idea comes from philosophical, cognitive, and computational considerations of how modality-specific information can be combined to form semantic representations. Wittgenstein (8) noted that, whereas experience was a key element to word meaning, there was a complex “criss-cross of partial feature similarities that come and go” when one considers the relationship between examples even of the same concept. Although some exemplars share enough features in common that they can be grouped on this basis, there are some exemplars that do not share any features in common yet belong to the same concept (e.g., games), or others where there is an open-ended structure to the partial similarities (e.g., handles). When some of the computational challenges are considered (detailed in the Introduction), it becomes readily apparent that concepts require more than the simple summation of semantic features (5–7, 9). There is a complex relationship between each

piece of modality-specific information and the concepts that it is associated with (Fig. 2). This complexity includes the facts that: features extend across a variable range of concepts; features do not reliably co-occur with each other either over time or over examples; features vary in conceptual diagnosticity (some features are superfluous, some are partial, and some are central to concepts); feature combinations are only a partial guide to concept identity (typical exemplars share many features, whereas atypical exemplars are those with a minimal feature overlap) (6, 28); and, certain feature combinations can change over time or relate to entirely novel exemplars, yet we are able to generalize our knowledge flexibly and appropriately to these new or changing items.

These challenges can be met if all of the modality-specific regions (the spokes in Fig. 1*B*) interact via a shared, modality-invariant representational system (the hub) (9). Because it is strongly cross-connected with multiple frontal, parietal, and temporal regions, the ATL provide a natural substrate to underpin the functioning of this representational hub. Using a semantic matching-to-sample paradigm, we were able to demonstrate that as the ATL are subjected to increasing atrophy—as it is in semantic dementia—semantic performance becomes increasingly dominated by superficial similarities. When patients were asked to select exemplars of the same concept from an array of pictures, they made a combination of errors. When nontarget items shared many features with the target concept, the patients had a strong tendency to choose these items; that is to say, they overgeneralized the concept. They also undergeneralized the same concepts in that they failed to select target exemplars. This tendency was especially pronounced when the target items were atypical. By combining a variety of targets and foils in the same task, we were able to demonstrate that the SD patients made under- and overgeneralization errors simultaneously. This is consistent with the notion that as the ATL hub representation breaks down, the category boundary not only becomes less certain but is also shifted and simplified.

How does this modality-invariant representation system work and how does it break down? By using an additional layer of representational units, computational models of semantic memory form a multidimensional space. Through a gradual training process, each piece of sensory, verbal, and motor information becomes associated with a subregion of this space. The closest analogy might be different kinds of geographical maps. Each type of map (e.g., geological, political, linguistic, agricultural, etc.) codes the same chart/grid system with the presence or absence of each type of feature (e.g., mountainous regions, wheat-growing areas, etc.) that is found in that modality (type of map). This shared representational space (the grid system used in all maps) results in a multilayered tapestry of information (shown in a simplified way in Fig. 2). Any specific location is then associated with (can reproduce) the information mapped to it (the name of the area, its geology, etc.), plus it can generate the likely information for areas that have never been directly mapped through interpolation (generalization to new examples).

By using many dimensions, it is possible to map complex non-linear regions—that is, map relationships between each feature and its associated concepts. In Fig. 4, we have shown the relationship between the verbal feature “cat” and a variety of different items. The items are arranged approximately by superficial visual similarity (coded perhaps in visual association cortex). To map “cat” correctly in this region via the ATL invariant representations, it is necessary to have a complex boundary (projected onto the items shown in Fig. 4*A*). As the representational hub breaks down, two things are likely to happen: The boundary becomes fuzzy, and is simplified because there are fewer dimensions to code the boundary in. Gradually, the boundary becomes increasingly dominated by superficial similarities (with reduced dimensions this is an optimal solution in that classification accuracy can be maximized by aligning the boundary with the shared feature structure). The consequence, however, is that two types of mismapping occur

(undergeneralizations) and false positive identification of the pseudotypical foils (overgeneralizations). The second change to standard word-picture matching was demanded by the need to observe over- and undergeneralizations separately. In standard word-picture matching there is only one target among a number of foils. When a patient fails the trial, they select one of the foils rather than the target. In this situation it is impossible to know whether this reflects an overgeneralization (the foil was considered to be the best example of the probe) or an undergeneralization (the target did not register as an example of the probe and thus the patient picks from one of the remaining items). We overcame this problem by placing multiple targets in each trial, asking the patients to pick which of the assembled pictures was an example of the probe and instructing them that there was always more than one target in the array. When the test is constructed in this way, it is possible to separate different kinds of error and to observe them simultaneously (*Results*).

The specific details of the test are as follows: The task consisted of 30 trials, all presented on a laptop. Each trial began by providing the patient with a written probe name (e.g., cat). Because SD patients sometimes fail to generate any information from a verbal input (32), we also provided a picture of the object (which was not repeated in the target array). The patients were then presented with an array of assorted pictures (an example is shown in Fig. S1) and were asked to pick which items were examples of the target item. The experimenter repeated the target name whenever requested by the patient. Having selected one or more items, the experimenter asked the participants whether there were any more examples of the item on the page. There were no time limits in the task. Each selection array contained nine pictures. There were always two or three targets in each array to prevent any strategies in terms of the number of targets per trial. The remaining items were foils. The

pictures can be divided into five different types. There were two types of target: typical targets and atypical targets. There were three types of foil: pseudotypical (items that shared many surface similarities with the target item); partially related (items that shared some of the surface similarities); and unrelated items. We included these different kinds of foil for two reasons. First, we used two types of related foil to test the hypothesis that there would be a graded effect in the emergence of overgeneralization errors—such that the mild patients would make these errors to the pseudotypical foils and that the more severe patients would also select the partially related items—reflecting a gradual loss of modality-invariant knowledge. The unrelated foils were included to check that the patients were not picking targets randomly, as even surface information with minimal invariant knowledge should be sufficient to reject these unrelated items (*Discussion*).

The data were tabulated according to the five types of item in each trial and the results are shown in Fig. 3. Performance in the patients was compared to that of seven age- and education-matched control participants. Patient scores were considered abnormal if they fell below the range of control performance on each separate condition.

ACKNOWLEDGMENTS. We are grateful to all the patients and the carers for taking part in this study. We thank Prof. Alistair Burns and Dr. Roland Zahn for referring some of the patients to us. E.J.M. was supported by an Overseas Research Scholarship and the University of Manchester. The research was supported by a programme grant from the Medical Research Council, UK (G0501632) and a cooperative grant from the Alzheimer's Research Trust (ART/NCG2005/1). M.A.L.R. held the M.R.C. grant funding. M.A.L.R. and R.W.J. held the A.R.T. funding.

- Eggert GH (1977) *Wernicke's Works on Aphasia: A Source-Book and Review* (Mouton, The Hague).
- Martin A (2007) The representation of object concepts in the brain. *Annu Rev Psychol* 58:25–45.
- Pulvermüller F (2005) Brain mechanisms linking language and action. *Nat Rev Neurosci* 6:576–582.
- Barsalou LW, Kyle Simmons W, Barbey AK, Wilson CD (2003) Grounding conceptual knowledge in modality-specific systems. *Trends Cogn Sci* 7:84–91.
- Murphy GL, Medin DL (1985) The role of theories in conceptual coherence. *Psychol Rev* 92:289–316.
- Smith EE, Medin DL (1981) *Categories and Concepts* (Harvard Univ Press, Cambridge, MA).
- Rogers TT, McClelland JL (2004) *Semantic Cognition: A Parallel Distributed Processing Approach* (MIT Press, Cambridge, MA).
- Wittgenstein L (2001) *Philosophical Investigations: The German Text, with a Revised English Translation 50th Anniversary Commemorative Edition* (Wiley-Blackwell, Oxford).
- Rogers TT, et al. (2004) Structure and deterioration of semantic memory: A neuropsychological and computational investigation. *Psychol Rev* 111:205–235.
- Plaut DC (2002) Graded modality-specific specialization in semantics: A computational account of optic aphasia. *Cogn Neuropsychol* 19:603–639.
- Catani M, Thiebaut de Schotten M (2008) A diffusion tensor imaging tractography atlas for virtual in vivo dissections. *Cortex* 44:1105–1132.
- Gloor P (1997) *The Temporal Lobe and the Limbic System* (Oxford Univ Press, Oxford).
- Morán MA, Mufson EJ, Mesulam MM (1987) Neural inputs into the temporopolar cortex of the rhesus monkey. *J Comp Neurol* 256:88–103.
- Wise RJ (2003) Language systems in normal and aphasic human subjects: Functional imaging studies and inferences from animal studies. *Br Med Bull* 65:95–119.
- Lambon Ralph MA, Patterson K (2008) Generalization and differentiation in semantic memory: Insights from semantic dementia. *Ann N Y Acad Sci* 1124:61–76.
- Patterson K, Nestor PJ, Rogers TT (2007) Where do you know what you know? The representation of semantic knowledge in the human brain. *Nat Rev Neurosci* 8: 976–987.
- Simmons WK, Martin A (2009) The anterior temporal lobes and the functional architecture of semantic memory. *J Int Neuropsychol Soc* 15:645–649.
- Bozeat S, Lambon Ralph MA, Patterson K, Garrard P, Hodges JR (2000) Non-verbal semantic impairment in semantic dementia. *Neuropsychologia* 38:1207–1215.
- Lambon Ralph MA, Lowe C, Rogers TT (2007) Neural basis of category-specific semantic deficits for living things: Evidence from semantic dementia, HSVE and a neural network model. *Brain* 130:1127–1137.
- Lambon Ralph MA, Pobric G, Jefferies E (2009) Conceptual knowledge is underpinned by the temporal pole bilaterally: Convergent evidence from rTMS. *Cereb Cortex* 19: 832–838.
- Pobric GG, Jefferies E, Ralph MA (2007) Anterior temporal lobes mediate semantic representation: Mimicking semantic dementia by using rTMS in normal participants. *Proc Natl Acad Sci USA* 104:20137–20141.
- Marinkovic K, et al. (2003) Spatiotemporal dynamics of modality-specific and supramodal word processing. *Neuron* 38:487–497.
- Vandenberghe R, Price C, Wise R, Josephs O, Frackowiak RSJ (1996) Functional anatomy of a common semantic system for words and pictures. *Nature* 383:254–256.
- Hodges JR, Patterson K, Oxbury S, Funnell E (1992) Semantic dementia. Progressive fluent aphasia with temporal lobe atrophy. *Brain* 115:1783–1806.
- Brambati SM, et al. (2009) Atrophy progression in semantic dementia with asymmetric temporal involvement: A tensor-based morphometry study. *Neurobiol Aging* 30:103–111.
- Nestor PJ, Fryer TD, Hodges JR (2006) Declarative memory impairments in Alzheimer's disease and semantic dementia. *Neuroimage* 30:1010–1020.
- Warrington EK (1975) The selective impairment of semantic memory. *Q J Exp Psychol* 27:635–657.
- Rosch E (1975) Cognitive representations of semantic categories. *J Exp Psychol Gen* 104:192–233.
- Lambon Ralph MA, McClelland JL, Patterson K, Galton CJ, Hodges JR (2001) No right to speak? The relationship between object naming and semantic impairment: Neuropsychological evidence and a computational model. *J Cogn Neurosci* 13: 341–356.
- Wooliams AM, Cooper-Pye E, Hodges JR, Patterson K (2008) Anomia: A doubly typical signature of semantic dementia. *Neuropsychologia* 46:2503–2514.
- Plaut DC, Shallice T (1993) Deep dyslexia: A case study of connectionist neuropsychology. *Cogn Neuropsychol* 10:377–500.
- Lambon Ralph MA, Howard D (2000) Gogi aphasia or semantic dementia? Simulating and assessing poor verbal comprehension in a case of progressive fluent aphasia. *Cogn Neuropsychol* 17:437–465.