UC Merced
Proceedings of the Annual Meeting of the Cognitive Science Society

Title
A PDP Simulation of the Effects of Transcranial Magnetic Stimulation on Semantic Cognition

Permalink
https://escholarship.org/uc/item/0mk3w7gp

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 31(31)

ISSN
1069-7977

Authors
Drakesmith, Mark
Pobric, Gorana
Welbourne, Stephen

Publication Date
2009

Peer reviewed

eScholarship.org Powered by the California Digital Library
University of California

A PDP Simulation of the Effects of Transcranial Magnetic Stimulation on Semantic Cognition

Mark Drakesmith (mark.drakesmith@postgrad.manchester.ac.uk)
Gorana Pobric (gorana.pobric@manchester.ac.uk)
Stephen Welbourne (stephen.r.welbourne@manchester.ac.uk)

Neuroscience and Aphasia Research Unit (NARU), School of Psychological Sciences, Zochonis Building, Brunswick Street, University of Manchester, Manchester, M13 9PL, UK.

Abstract

Transcranial magnetic stimulation (TMS) is a useful tool for determining cortical interactions that take place during semantic cognition (Pobric *et al*, 2007). TMS of the anterior temporal lobe (ATL) induces specificity-graded conceptual breakdown of stored representations resulting in differential effects on naming tasks depending on the level of specificity required. This paper aims to elucidate this effect using the Rogers *et al* (2004) model of semantic cognition. TMS effects are modelled by reducing the gain of the affected units to simulate reduction in neuronal sensitivity. Results show that basic-level naming is more robust to rTMS than subordinate level naming as found by Pobric *et al*. In addition the model predicts that superordinate naming should be even more robust to rTMS than basic level naming. This specificity-graded breakdown of semantic memory appears to be independent of learnt word frequency. This supports evidence that the ATL's function is that of category differentiation.

Keywords: neuroscience; psychology; cognitive architectures; concepts and categories; distributed cognition; memory; semantics; computational neuroscience; computer simulation; neural networks

Introduction

Semantic memory has become a key area for neuropsychological and neuroimaging research. There are a number of characteristic semantic effects such as typicality, context, and familiarity, which can be used to probe the organisation of semantic representations.

One characteristic feature of semantic memory is that when naming an item, there is strong bias towards a 'basic' level of categorisation. For example when asked to name a picture of a brown Labrador, the most common response is 'dog', rather than a subordinate name, like 'Labrador' or a superordinate name such as 'animal' (see Murphy & Lassaline, 1997 for a review). A proposed explanation of the basic-level effect is that of category differentiation, which arises from a balance between informativeness and distinctiveness (Murphy & Brownell, 1985). Subordinate names are the most informative while superordinate names are the most distinctive.

Neuropsychological and Neuroimaging Findings Much insight into the neural basis of semantic cognition has been derived from patients with Semantic dementia (SD) which in particular occurs in the anterior temporal lobes (ATL) and almost exclusively results in impaired semantic memory (see Hodges & Patterson, 2007, for a review). One feature common to many SD patients is poor naming performance, which is differentiated according to naming level: subordinate naming is the worst affected whilst superordinate naming is the most likely to be preserved.

Evidence from neuroimaging studies such as PET and fMRI reveal two contradictory perspectives on how semantic memory is stored. Some (e.g. Petersen *et al*, 1988; Grabowski *et al*, 2001) have shown that participants given visual and verbal stimuli for various items, show increased blood flow to frontal and temporal regions of the cortex. However, there is also evidence that semantic memory is distributed. Thompson-Schill *et al* (1999) found that participants asked to recall visual properties of an object showed increased activity in the left fusiform gyrus. Grosman *et al* (2002) found that retrieval of actions associated with tool use elicited higher activation of pre-motor cortex, suggesting a store of motor representations. This distributed theory has been used to explain why there appear to be distant neural substrates for animal objects and man-made objects (e.g. Caramazza & Shelton, 1998) because of differing perceptual content between these categories.

These two contradictory strands have been unified in the form of the semantic hub theory (reviewed in Patterson *et al*, 2007). The theory states that semantic memory is distributed across regions which store representations specific to a certain modality such as visual or motor representations. These representations are, however, interconnected via a central 'semantic hub' in the ATL. This hub gives rise to amodal semantic concepts through its interconnections between the modality-specific representations.

Some imaging studies have looked at naming level effects (e.g. Tyler *et al*, 2004) and have suggested a continuum of naming specificity along the posterior-anterior axis of the inferior temporal cortex (Martin & Chao, 2001). However, a meta-analysis by Patterson *et al* (2007) shows subordinate name activations are usually no more anterior than those for basic-level naming. They suggested alternatively that subordinate naming requires activation of more anterior regions of temporal cortex because of the increased computational demands of differentiation of overlapping concepts within a particular category. This argument parallels with the informativeness-distinctiveness account for naming level effects (Murphy and Brownell, 1985). It is also consistent with that of Rogers *et al* (2007) (described below) which argues that naming-level effects are an emergent property of the semantic network.

Transcranial Magnetic Stimulation (TMS) TMS has been used as an investigative technique in cognitive neuroscience since the 1980s (reviewed in Walsh & Pascual-Leone, 2003). Unlike passive imaging techniques, TMS can be used to alter the brain's dynamics and thus establish causal links between brain function and cognition.

The TMS paradigm typically used in cognitive neuroscience is a repeated constant train of TMS pulses (rTMS) usually at 1Hz for approx. 5-10 minutes. Participants are usually tested on a behavioural task before and after TMS. Effects of TMS manifest in changes in reaction times (RTs) and rarely in obvious behavioural changes.

rTMS provides new avenues to explore the neural basis of semantic memory. Pobric *et al* (2007) applied rTMS to the ATL and measured RTs for a naming task. Participants were asked to produce basic level or subordinate level names to visual images. The results show a significant effect of rTMS on subordinate naming but not on basic naming.

This finding adds further evidence to indicate the ATL is required to differentiate overlapping semantic concepts. Although the extent of the effect of rTMS is much smaller (effects only manifest in RTs), the mechanisms of conceptual breakdown are considered analogous.

Physiology of TMS There is still debate of how TMS exerts its disruptive effect. One account is that it introduces aberrant noise into the neural dynamics (Walsh & Coway, 2000). This does not account for the relative persistence of TMS effects with no further stimulation, or the robustness of neural populations to noisy inputs.

Recent evidence suggests synaptic plasticity as the cause (see Thickbroom, 2007 for a review). Some studies find the suppressive or facilitatory TMS effects are rate-dependant (Fitzgerald *et al*, 2006). High frequency pulse trains (about 5Hz or over) have LTP-like effects, while low frequency pulse trains, as commonly used for investigation of cognitive function, result in LTD-like effects.

Huang *et al* (2007) tested the effects of TMS on LTP/D by administering participants with an NMDA antagonist; NMDA being a vital precursor to LTP induction. They found low-frequency TMS applied to participants given the NMDA antagonist failed to suppress evoked motor responses compared those given placebos. Also, Esser *et al* (2006) measured ERPs before and after high-frequency rTMS. They found amplitudes of ERPs recorded near the stimulation site were significantly higher after rTMS, indicating an LTP-like effect. The reductions in RT in behavioural experiments have therefore been attributed to reduced sensitivity of the cortex to synaptic inputs. Some studies of TMS on motor cortex have measured changes in the input-output (I/O) function (e.g. Muellbacher *et al*, 2000). This is a convenient analogy to make with sigmoid activation function used in many PDP models. Reduction of synaptic sensitivity can be simulated by reducing the gain of this function.

Computational Models Parallel distributed processing (PDP; Rumelhart *et al*, 1986) is a popular approach to studying cognition because of its ability to produce models that learn and have biologically inspired features. There has been much interest in modelling semantic memory in a PDP framework, and particularly, in lesioning such models to simulate symptoms of neurological conditions such as SD (reviewed in McClelland & Rogers, 2003),

Many of these models have simply tried to model the behavioural characteristics of semantic memory and have, until recently, remained divorced from the underlying cortical structure. In the spirit of the semantic hub theory, Rogers *et al* (2004) proposed a model consisting of separate layers corresponding to different modality-specific representations in different cortical regions. All of these regions are linked via the semantic layer (representing the ATL) which mediates links between modality-specific features.

The model was trained on items derived from a set of prototypes, themselves derived from a hierarchical cluster analysis of semantic norms (Garard *et al*, 2001). The model was trained on both subordinate and basic names. However, it was not specifically tested for naming level effects, although a range of deficits that correlate with SD were demonstrated.

Although level effects were not explicitly tested in this model, Rogers *et al* (2007) subsequently made predictions about how such effect would arise because of movement through the semantic representational space. These predictions match Murphy and Brownell's (1985) distinctiveness-informativeness account. As the network trains, items rearrange into clusters in the representational space. These clusters reflect similarities between items: Items that are very similar organise into dense, tightly packed clusters. Intra-cluster density is high. Therefore, concepts within clusters (at the subordinate level) are informative but not distinctive. Inter-cluster density is low and therefore concepts between clusters (at the superordinate level) are distinctive but uninformative. This type of organisation was shown in a previous model by Rumelhart & Todd (1993) and by Rogers & McClelland (2004) but has not yet been tested in the multi-modal network.

Aims The aim of this project is to simulate the effects of TMS found by Pobric *et al* (2007) using the Rogers *et al* (2004) model. As well as showing behavioural aspects of normal and impaired semantic cognition, the model's multi-modal architecture is capable of replicating the experimental paradigm of Pobric *et al* (2007). The model will be used to determine if normal naming level effects are emergent property of the network or dependant on word frequency (Brown, 1958). The model can then be 'lesioned', but rather than removing connection weights, rTMS is simulated by reducing the gain of the sigmoid activation function within the semantic layer. This is aimed to simulate the reduction in synaptic sensitivity in this region as shown by physiological evidence.

Method

All network simulations are carried out in LENS neural network simulator (Rohde, 2003) with subsequent analysis carried out in MATLAB. The model is adapted from that of Rogers *et al* (2004). Two frequency manipulations were made (flat vs. basic-level inflated frequencies). All layers are connected via bidirectional, dense, asymmetrically weighted links. Three main type of representation are given: visual features, verbal descriptions and names. The visual features are encoded in one layer. The verbal descriptions are encoded in three layers, subdivided into perceptual, functional and encyclopaedic descriptions.

The naming layer is expanded slightly from the original model. (see figure 1). Each name is encoded by the activation of one unique unit in each layer. During training, the model is presented with separate examples for each level of naming and for each modality. Targets for each modality are given during training and the name appropriate for the specified level is also given. A name is generated in all three levels and not restricted to a particular level. This is equivalent to the free-recall paradigm in experimental psychology.

The network is trained on 48 items, each associated with unique subordinate names, 6 basic names (8 items for each name) and 2 superordinate names (24 items for each name).

All units in the network have a dot product input function and a sigmoid output function. The input is also subject to an untrainable bias, which is set to -2 throughout the network. The network simulates 'continuous' time over seven time intervals; each time interval is further divided into 10 ticks, which are arbitrary units of 'continuous' time. The input function is treated as a derivative of each units input. To obtain the actual input for each tick, the Euler method is used to integrate the input function over time.

The weights are initially set at random, and then during training are modified by a back-propagation algorithm using the delta rule. While training, the network is presented with inputs of each item in each modality by clamping the relevant layers to the relevant stimuli. The network is then permitted to cycle for 3 intervals. Inputs are then removed for 2 more intervals and then targets applied for 2 more intervals. Since there are 10 ticks per interval, the maximum possible length of each example is 70 ticks. Error derivatives are only calculated for each unit on each tick for which the network is given a target. Training is also subject to a learning rate of 0.005 and a weight decay of 0.0002. The total training length for the network is 10,000 epochs.

The training examples were selected in a probabilistic manner, with the constraint that each name should be presented equally often; because the number of names in each level does not match the number of items, the frequency of each name level has to be modified to satisfy this constraint. For example, there are 48/2=24 examples presenting the same superordinate name in the training set and each basic name will be 48/6=8 times more likely appear in the training set. The frequency of each example is therefore matched to the number of names in that examples' naming level. So f_{sub}=48, f_{bas}=6 and f_{sup}=2. The first experiment was carried out using the flat name frequencies as above. The second experiment used an artificially inflated basic-level frequency (f_{bas}=60) to simulate the fact that basic level names tend to be higher frequency than subordinate level ones.

Once trained, the network is tested on each item by clamping the relevant group for 3 cycles. The RT is defined as the tick at which all units are within a threshold of 0.3 of their target activations. i.e. the RT is when all ON-target units have activations above 0.7 and all OFF-target units have activations below 0.3. This level was found to be

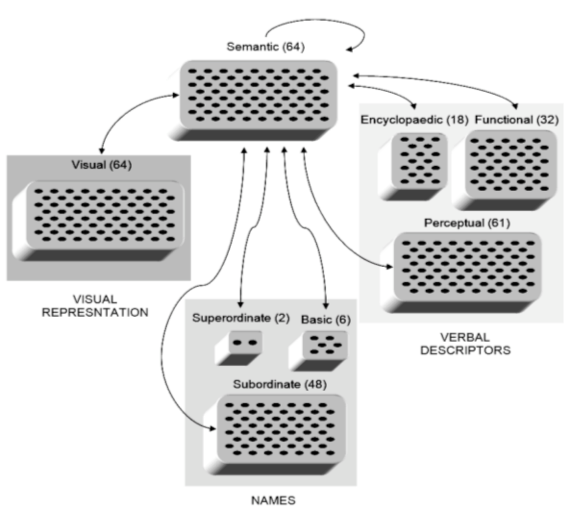

Figure 1: Structure of the PDP network for the non-competitive naming simulations. Numbers in brackets indicate number of units.

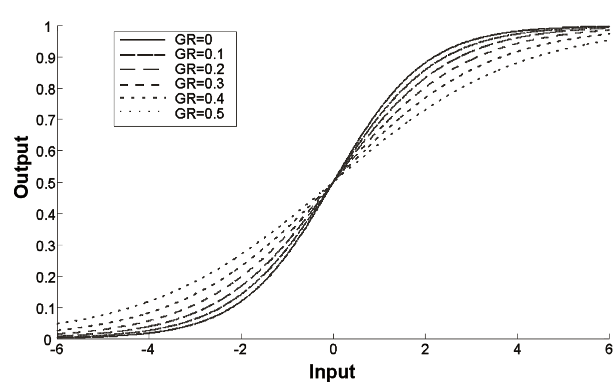

Figure 2: Sigmoid activation function of semantic units (GR=0) and the gain reduction manipulations used to simulate rTMS.

rigorous enough to prevent potentially ambiguous patterns of activation being reported as decisive responses, but also lenient enough to allow naming to be at or near 100% accuracy.

The effects of rTMS were simulated in the PDP environment by altering the gain of the sigmoid activation function to mimic the effect of a reduction in synaptic sensitivity. Gain reductions of 0 to 0.5 in increments of 0.1 were tested (figure 2).

Results

Significance tests between different levels of naming were carried out with Welch's *t*-tests.

With the flat-frequency network, RTs between all levels are significantly different (figure 3a), with subordinates having the highest RT and superordinates having the lowest RT. This pattern is reflected in the mean unit activation during trials, with superordinate name activations reaching threshold first, followed by basic, and then subordinate naming (figure 3b).

The effects of TMS (as simulated by gain reduction) reveal a specificity-graded breakdown of naming that is congruent with the results of Pobric *et al*'s behavioural data. Subordinate naming shows the strongest TMS effect, followed by basic naming, with superordinate names showing the weakest TMS effect. The TMS effect for subordinate naming becomes significantly higher than for the other naming levels at a gain reduction of 0.3. The difference in TMS effect between basic naming and superordinate naming only becomes significant at a gain reduction of 0.5 (figure 3c).

With the basic level inflated frequencies, initial RTs for basic naming are significantly lower than for both subordinate and superordinate naming, as in human naming. There is no significant difference between subordinate and superordinate naming (figure 4a). The mean unit activation for basic names reaches threshold much sooner than subordinate and superordinate naming, the rate of ascent of the basic name activations being higher in the basic-inflated network than the flat-frequency network (figure 4b).

The effects of gain reduction on RTs are similar to the flat-frequency network and to the results from Pobric *et al.* The TMS effect for subordinate naming becomes significantly higher than for the other levels at a gain reduction of 0.3. Meanwhile basic naming, although showing a similar TMS effect never becomes significant

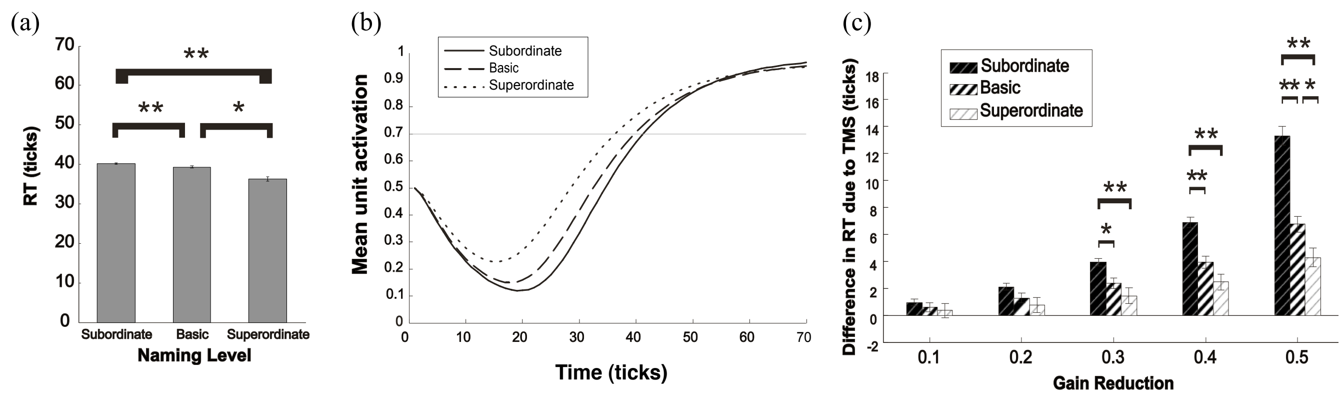

Figure 3: Results of simulations on network trained on flat frequencies (a) RTs for naming a different levels (b) Time course of target-ON name unit activations, averaged over all examples. (c) Effects of gain reduction on RT for each naming level. In (a) and (c) error bars indicate standard error and significant differences between naming levels are indicated by braces: $p<0.05$ indicated by * and $p<0.001$ indicated by **.

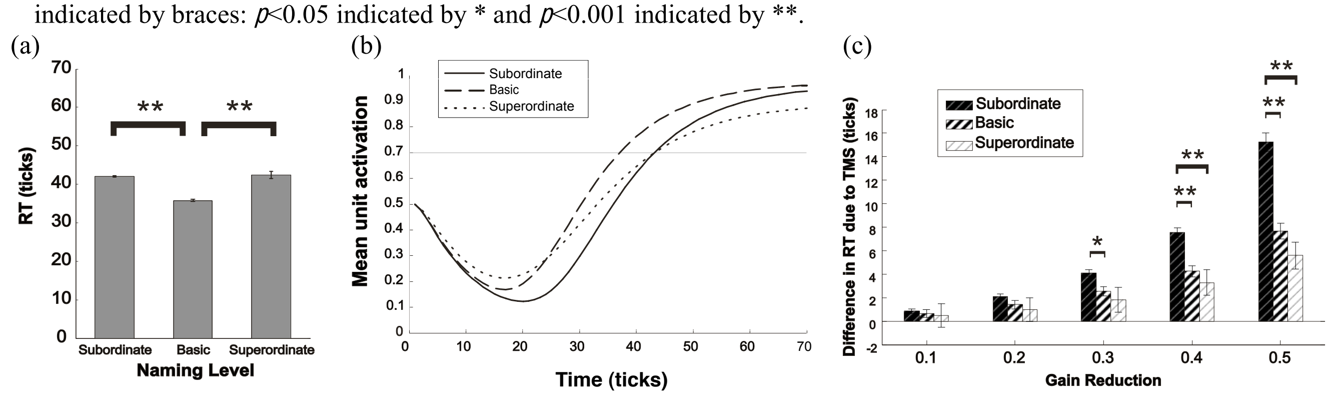

Figure 4: Results of simulations on network trained on basic-level inflated frequencies. (a) RTs for naming a different levels (b) Time course of target-ON name unit activations, averaged over all examples. (c) Effects of gain reduction on RT for each naming level. In (a) and (c) error bars indicate standard error and significant differences between naming levels are indicated by braces: $p<0.05$ indicated by * and $p<0.001$ indicated by **.

compared to superordinates (figure 4c). Nevertheless there is still a persistent trend for superordinates to be more robust to TMS effects and it appears it would continue through greater gain reductions, which is likely to yield a significant TMS effect for basic names. The robustness of superordinates to TMS in the basic-inflated network is surprising as the overall RT for basic naming is much lower then for superordinate naming (figure 4c) suggesting that it is an easier task and so should be more resistant to the effect of rTMS than superordinates.

Discussion

This project aimed to model impairments in semantic processing as a consequence of rTMS to ATL (Pobric *et al*, 2007) using a pre-existing PDP framework (Rogers *et al*, 2004). Results show that the network with inflated basic-name frequencies gave rise to both basic-level superiority in undamaged naming and selective impairment of subordinate, but not basic level, naming by rTMS. These results are in accordance with the human data from the study by Pobric *et al* (2007). That study did not obtain results for superordinate naming. However, the model predicts that superordinate naming should be more stable than basic-level naming and so will be unimpaired by rTMS at the dosage levels used in the Pobric *et al* study. In some ways this is a surprising prediction as superordinate naming is slower than basic level naming which might be taken as a sign that it is a harder task and so should be more susceptible to TMS. However, it is consistent with the neuropsychological evidence of SD patients who show impaired subordinate naming but relatively unpaired superordinate naming. This suggests that the analogy previously made between rTMS and SD is valid.

The flat-frequency condition shows that the emergent naming level effects in the model do not fully account for the human naming pattern, although Rogers *et al* (2007) predicted that they would. Only when basic-level name frequencies are inflated in the second network, does basic-level naming superiority arise. This suggests that basic-level superiority is only partially an emergent property of the network and also depends on frequency.

Although human-like naming-level effects did not manifest themselves without frequency manipulation, there is still a specificity-graded breakdown of conceptual knowledge caused by rTMS in both networks (figure 3c & 4c). This suggests that the specificity-graded effects of rTMS are independent of name frequency. These results add support for Patterson *et al*'s (2007) proposition that the ATL's role is not so much in the storage of highly specific information, but in the differentiation of highly overlapping concepts. It also supports Rogers *et al* (2007) position that such an ability is emergent.

Further information about the structure of semantic representations can be ascertained by examining the behaviour of the semantic layer during testing. Dimensional scaling and cluster analysis of the learnt semantic representations can reveal patterns reflecting naming specificity. If Roger's *et al* (2007) predictions hold true, we can expect to find a pattern of subordinate categories organized in small and dense (informative but indistinctive) clusters while superordinate categories being in large and sparse (distinctive but uninformative) clusters. As a result, basic level categories would show high, intra-cluster density (informative) and low inter-cluster density (distinctive).

Based on the speculated cluster structure, Rogers *et al* (2007) made predictions about the rate of unit activations: unit activation of target-ON units are dictated by the network's movement through representational space as it tended towards the target item. Initially, superordinate items were predicted to activate first, followed by basic and finally subordinate names. However, the dense clustering of basic and subordinate names results in subsequent rapid activations for those levels, and consequently, in basic names reaching threshold first. The mean unit activations in the basic-inflated network (figure 4b) shows this predicted pattern of activation. The rates of initial descent correspond to the naming levels, with subordinates showing the sharpest decent. Subordinates and basics subsequently show faster ascents than superordinate, but because basic names had a slower initial decent, it reaches threshold before either subordinate or superordinate.

Again, this pattern was not observed in the flat-frequency network, suggesting again that this effect is to some extent frequency-dependant. However, the differing rates of ascent for superordinates and subordinates in this network indicate only a partial dependence on frequency. The structure of the learnt semantic representations also clearly plays a role.

It should be noted that Roger *et al*'s (2007) prediction was made for semantic unit activations rather than name units. The decision to make a response in a semantic task is likely to involve a network comprising temporal lobe and pre-frontal cortex (e.g. Jennings *et al*, 1997) and not in cortical regions responsible for modality-specific representations. The name units themselves only activate once threshold is reached in the semantic units so one would expect them to display similar time courses.

Another issue with the model that should be considered is that there is no delay propagation between units. Although units activate in a temporally graded manner, the spreading of activations to other linked units do not. When visible units are clamped to ON, the activation of the connected semantic units progress in a graded manner, but then all other visible units connected via those semantic units also activate simultaneously. Propagation delays between cortical populations during cognitive tasks, often connected via polysynaptic connections, are likely to significantly contribute to RTs.

Conclusion The basic level superiority effect found in normal naming and the differential inhibitory effects of rTMS on naming demonstrated by Pobric *et al* (2007) have been successfully replicated in a PDP model based on the semantic hub theory. The basic-level superiority effect is largely influenced by increased frequency as previously indicated (e.g. Brown, 1958). However, there is also indication that internal structure of the learned semantic representations, rather than frequency, is responsible for the differential rTMS effect on naming. The model also predicts that rTMS should have a greater effect on basic level naming than on superordinate naming. Although this appears counter intuitive, based on superordinate naming being an apparently harder task, it is consistent with the idea that there is a central inter-modal hub responsible for category differentiation.

References

Brown, R. (1958). How shall a thing be called? *Psychological Review, 65*, 14-21.

Caramazza, A., & Shelton, J. R. (1998). Domain-specific knowledge systems in the brain the animate-inanimate distinction. *Journal of Cognitive Neuroscience, 10*, 1–34.

Castellani, G. C., Quinlan, E. M., Bersani, F., Cooper, L. N., & Shouval, H. Z. (2005). A model of bidirectional synaptic plasticity: From signalling network to channel conductance. *Learning and Memory, 12*, 423-432.

Esser, S. K., Huber, R., Massimini, M., Peterson, M. J., Ferrarelli, F., & Tononi, G. (2006). A direct demonstration of cortical LTP in humans: A combined TMS/EEG study. *Brain Research Bulletin, 69*, 86-94.

Fitzgerald, P. B., Fountain, S., & Daskalakis, Z. J. (2006). A comprehensive review of the effects of rTMS on motor cortical excitability and inhibition. *Clinical Neurophysiology, 117*, 2584–2596.

Grabowski, T. J., Damasio, H., Tranel, D., Boles Ponto, L. L., Hichwa, R. D., & Damasio, A. R. (2001). A role for left temporal pole in the retrieval of words for unique entities. *Human Brain Mapping, 13*, 199-212.

Huang, Y.-Z., Chen, R.-C., Rothwell, J. C., & Wen, H.-Y. (2007). The after-effect of human theta burst stimulation is NMDA receptor dependent. *Clinical Neurophysiology, 118*, 1028–1032.

Jennings, J. M., McIntosh, A. R., Shitij Kapur, S., Tulving, E., & Houle, S. (1997). Cognitive subtractions may not add up: The interaction between semantic processing and response mode. *NeuroImage, 5*, 229-239.

Kucera, H., & Francis, W. N. (1967). *Computational Analysis of Present-day American Engish*. Providence: Brown University Press.

Martin, A., & Chao, L. L. (2001). Semantic memory and the brain: structure and processes. *Current Opinion in Neurobiology, 11*, 194-201.

McClelland, J. L., & Rogers, T. T. (2003). The parallel distributed processing approach to semantic cognition. *Nature Reviews Neuroscience, 4*, 310-322.

Muellbacher, W., Ziemann, U., Boroojerdi, B. & Hallett, M. (2000). Effects of low-frequency transcranial magnetic stimulation on motor excitability and basic motor behavior. *Clinical Neurophysiology, 111*, 1002-1007.

Murphy, G. L., & Brownell, H. H. (1985). Category differentiation in object recognition: typicality constraints on the basic category advantage. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 11*, 70-84.

Murphy, G. L., & Lassaline, M. E. (1997). Hierarchical Structure in Concepts and the Basic Level of Categorization. In K. Lamberts & D. R. Shanks (Eds.), *Knowledge, Concepts and Categories* (pp. 93-131). Hove: Psychology Press.

Patterson, K., Nestor, P. J., & Rogers, T. T. (2007). Where do you know what you know? The representation of semantic knowledge in the human brain. *Nature Reviews Neuroscience, 8*, 976-987.

Petersen, S. E., Fox, P. T., Posner, M. I., Mintun, M., & Raichle, M. E. (1988). Positron emission tomographic studies of the cortical anatomy of single-word processing. *Nature, 331*, 585-589.

Pobric, G., Jefferies, E., & Lambon Ralph, M. A. (2007). Anterior temporal lobes mediate semantic representation: Mimicking semantic dementia by using rTMS in normal participants. *Proceedings of the National Academy of Sciences, 104*, 20137–20141.

Rohde, D. (2003). LENS: The light, efficient network simulator (Version 2.63).

Rogers, T. T., Lambon Ralph, M. A., Garrard, P., Bozeat, S., McClelland, J. L., Hodges, J. R., Patterson, K. (2004). Structure and deterioration of semantic memory: A neuropsychological and computational investigation. *Psychological Review, 111*, 205–235.

Rogers, T. T., & McClelland, J. L. (2004). *Semantic Cognition: A Parallel Distributed Processing Approach*. Cambridge: MIT Press.

Rogers, T. T., & Patterson, K. (2007). Object categorization: Reversals and explanations of the basic-level advantage. *Journal of Experimental Psychology: General, 136*, 451–469.

Rumelhart, D. E., McClelland, J. L., & The PDP Research Group. (1986). *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*. Cambridge: MIT Press.

Thickbroom, G. W. (2007). Transcranial magnetic stimulation and synaptic plasticity: Experimental framework and human models. *Experimental Brain Research, 180*, 583–593.

Thompson-Schill, S. L. (2003). Neuroimaging studies of semantic memory. *Neuropsychologia, 41*, 280–292.

Thompson-Schill, S. L., Aguirre, G. K., D'Esposito, M., & Farah, M. J. (1999). A neural basis for category and modality specificity of semantic knowledge. *Neuropsychologia, 37*, 671–676.

Tyler, L. K., Stamatakis, E. A., Bright, P., Acres, K., Abdallah, S., Rodd, J. M., Moss, H. E. (2004). Processing objects at different levels of specificity. Journal of Cognitive Neuroscience, 16, 351-362.

Walsh, V., & Cowey, A. (2000). Transcranial magnetic stimulation and cognitive neuroscience. *Nature Reviews Neuroscience, 1*, 73–79.

Walsh, V., & Pascual-Leone, A. (2003). *Transcranial Magnetic Stimulation: Towards a Neurochronometrics of Mind*. Cambridge: MIT Press.