

A Tale of Two Semantic Systems: Taxonomic and Thematic Knowledge

Daniel Mirman (mirmand@einstein.edu)
Grant M. Walker (walkerg1@einstein.edu)
Kristen M. Graziano (graziank@einstein.edu)

Moss Rehabilitation Research Institute
50 Township Line Rd., Elkins Park, PA 19027, USA

Abstract

Behavioral, neuroimaging, and lesion analysis data suggest two parallel semantic systems. One system, with anterior temporal lobe as critical hub, captures taxonomic relations based on feature overlap. A second system, with temporo-parietal junction as critical hub, captures thematic relations based on complementary roles in events. We describe a computational model of this theory that accounted for a one-way behavioral dissociation in aphasic picture naming errors (more taxonomic errors than thematic errors) and a neuroanatomical double dissociation (damaging feature representations led to relatively more taxonomic errors, damaging event representations led to relatively more thematic errors). The model also predicted that both taxonomic and thematic competitors should be automatically activated during single word processing, with taxonomic competitors activated more quickly and more strongly. These predictions were tested and confirmed in a spoken word comprehension experiment using eye tracking to assess the time course of competitor activation.

Keywords: semantic knowledge; taxonomic relations; thematic relations; event processing; computational modeling; spoken word processing.

Introduction

A core question in cognitive science is how semantic knowledge is represented. The study of semantic knowledge is typically intertwined with the study of feature-based or hierarchical conceptual categories. Feature-based accounts can explain a very broad range of phenomena (e.g., Rogers & McClelland, 2004) and are particularly effective in capturing the categorical, or taxonomic, structure of conceptual knowledge (e.g., Rogers & McClelland, 2004; O'Connor, Cree, & McRae, 2009). However, thematic conceptual knowledge – the grouping of concepts by participation in the same scenario or event (e.g., Estes, Galonka, & Jones, 2011) – is not as well captured by traditional feature-based accounts. On feature-based accounts, semantic similarity is a function of feature overlap (e.g., Cree, McRae, & McNorgan, 1999; Mirman & Magnuson, 2009; Rogers & McClelland, 2004), but thematically related objects typically share few, if any, features. Rather, they have complementary features that are related to the complementary roles the objects play in events.

There is a long history of behavioral studies demonstrating that thematic knowledge plays an important role in adult conceptual knowledge (e.g., Goldwater,

Markman, & Stilwell, 2011; Hare et al., 2009; Lin & Murphy, 2001; for a review, see Estes et al., 2011). One recent study used voxel-based lesion-symptom mapping (VLSM) to examine the neural basis of taxonomic and thematic processing (Schwartz et al., in press). Schwartz et al. analyzed picture-naming errors generated by a large sample of individuals with left hemisphere stroke aphasia (N=86), distinguishing between taxonomic errors (e.g., apple named as “pear” or “grape”) and thematic errors (e.g., apple named as “worm” or dog named as “bone”). Taxonomic errors were defined as category coordinates, superordinates, or subordinates; thematic errors were defined as incorrect responses which come from a different category but frequently play a complementary role with the target in events. The behavioral results showed a single dissociation: there were far more taxonomic errors than thematic errors (approximately 5:1 ratio) and all but two patients made more taxonomic errors than thematic errors. However, the lesion analysis revealed a neuroanatomical double dissociation in the relative proportion of taxonomic to thematic errors. Patients with lesions affecting the left anterior temporal lobe (ATL; Brodmann area 38 and the anterior portions of BA 20 and 21) tended to produce a higher proportion of taxonomic errors relative to thematic errors. In contrast, patients with lesions affecting the left temporo-parietal junction (TPJ; BA 39, posterior BA 21 and 22, superior BA 37, and BA 41 and 42) tended to produce a higher proportion of thematic errors relative to taxonomic errors.

The ATL effect is consistent with previous studies demonstrating its critical role in lexical-semantic processing (e.g., Hodges, Graham, & Patterson, 1995; Lambon Ralph et al., 2001; Patterson, Nestor, & Rogers, 2007; Schwartz et al., 2009). The TPJ effect makes contact with studies suggesting an important role for TPJ in thematic relations (e.g., Kalenine et al., 2009) and relational knowledge more generally (e.g., Wu, Waller, & Chatterjee, 2007; for a recent comprehensive review of neuroimaging studies of semantic representations see Binder et al., 2009).

Our first goal was to develop a formal computational model of these complementary semantic systems that can account for the neuroanatomical double dissociation as well as the one-way behavioral dissociation. Our model is related to previous work by Plaut (1995), who distinguished between semantic *relatedness* based on semantic feature overlap and semantic *association* based on temporal co-

occurrence to account for differences between associative and semantic priming effects in visual word recognition. We tested whether this distinction, combined with an explicit event representation, can capture both the one-way behavioral dissociation and the neuroanatomical double dissociation.

Like previous models of related phenomena (e.g., Plaut, 1995), our model predicts that thematically-related concepts are automatically activated during single word processing even when such activation is not required by task demands. This prediction is consistent with recent priming studies that demonstrate fast activation of event-based relations (e.g., Hare et al., 2009). We further tested this prediction in a spoken word-to-picture matching task with eye tracking to examine the time course of taxonomic and thematic competition.

Computational Model

The model follows in the parallel distributed processing tradition of modeling cognition as the bidirectional, graded, and nonlinear interactions among many simple processing units. Each unit is associated with an activation state, which is determined by the strength of its input and a sigmoid activation function. Units interact through weighted connections and the weights are learned over the course of training. The structure of the model was based on three core principles: (1) Taxonomic structure: individual concepts are represented by sets of semantic features, which tend to be shared by concepts within a category. (2) Thematic/event structure: normal word production is situated in an event or sentence context, which imposes regularities on which objects will co-occur. (3) Distinct representation of event information and semantic feature information based on the neuroanatomical and psychological evidence.




Figure 1. Architecture of the model.

Model Architecture

The model had 4 groups of units representing Semantic Features, Events, Lexical Semantics, and the Output Lexicon. The architecture of the model is shown in Figure 1. The arrows indicate full connectivity between layers (and fully recurrent connections between units in the Lexical

Semantics layer). Since the model was primarily designed to capture word production data, external input was provided to the Semantic Features and Event layers. Conceptually, the model was trained to perform a simplified analog of event description: there was a constant event input (i.e., the model is describing a single, ongoing event), a sequence of individual concept representations chosen (input) for production, and a corresponding sequence of target word outputs. The model was tested on picture naming by presenting input to the Semantic Features layer and evaluating the Output Lexicon activations. The simulations were conducted in LENS: (www.stanford.edu/group/mvc/LENSManual/index.html).

Representations

The Output Lexicon was a localist representation of 16 words where each unit represented a unique concept name. The 16 concepts were divided into 4 categories, each with 4 category members. The 32 units in the Semantic Features layer were divided into 4 sets of 8 units, with each set representing the possible features for a single category. For each concept there were 2 category coordinates that shared 50% of the object's features and the remaining category coordinate shared 0% of the object's features. Across categories, concepts shared 0% of their features.

The 8 units in the Event group represented four general event types (e.g., eating) with two directions for each event that implicitly specify the roles of the two event participants (e.g., “eats” vs. “is eaten by”). Each event involved two participants from a set of four possible concepts (one from each category).

Model Training

The model was trained on a simplified analog of sentence production. All of the sentences consisted of two concept names produced in a particular sequence that coincided with semantic feature input that changed from concept to concept and constant event representation input. At the start of each training trial, activations for all units were initialized to a small value (0.2), with a default soft-clamp (clamp strength: 0.2) at a value of 0.0. When groups received external input, their clamp strength was changed to 0.8. First, one event unit was soft-clamped to a value of 1.0, then the semantic features corresponding to the first concept in the sentence were also soft-clamped to a value of 1.0. Once these inputs were set, activation was allowed to propagate through the network for up to 10 time ticks with target activations defined for the output lexicon and semantic features layers. This corresponded to the production of the first word in the sentence. At the end of those 10 ticks, or if all activations were within 0.5 of their targets, all targets were removed and new external inputs for the semantic features layer were set, corresponding to the second concept in the sentence. Activation was again allowed to propagate through the network for up to 10 ticks with target activations defined for the semantic features and output lexicon. The connection

weights were updated after each training trial (i.e., production of each two-word sentence) using the error back-propagation through time algorithm (Pearlmutter, 1995) with learning rate set to 0.1 and momentum set to 0.9.

There were 32 possible example sentences: 4 event types with 4 possible first concepts and 2 possible second concepts. The model was trained for 5000 trials, at which point cross-entropy error had approached an asymptotic minimum.

Model Testing

The critical model test was an analog of the picture naming task. External inputs specifying a single concept were hard-clamped to the semantic features units and no external input was provided to the event layer. Output lexicon activations were recorded for 20 ticks. No weight changes occurred on test trials. To compute a model analog of picture naming response, we considered the cumulative activation received by each word type (target, taxonomic competitor, and thematic competitor) over the first 15 ticks of processing (since at that point the target word was no longer the most active output unit). These summed activations were then normalized to compute a proportion of activation received by each response type. ATL lesions were modeled by removing 20% of the connections between the semantic features layer and the lexical semantics layer; TPJ lesions were modeled by removing 50% of the connections between the event layer and the lexical semantics layer.

Results and Discussion

The simulations were repeated for 20 models, using different random starting weights and different random lesions. Figure 2 shows the average activation patterns for the target word, its taxonomic competitors, and its thematic (event) competitors for the fully trained control (unlesioned) model (left panel), following a lesion disrupting communication between event representations and lexical semantics (i.e., a TPJ lesion, middle panel), and following a lesion disrupting communication between semantic features and lexical semantics (i.e., an ATL lesion, right panel). The

target word was by far the most active word, indicating that although the model was trained to perform two-word sentence production, it was quite capable of performing single word production without an event context (i.e., no event input). This was true even after damage, indicating that our lesion implementation did not eliminate naming performance. Critically, both taxonomic and thematic competitors were activated, though with somewhat different time courses. Taxonomic competitors were strongly activated early in concept processing, but this activation was transient, peaking well before the target activation reached its peak (for a similar time course in spoken word recognition, and its implications, see Mirman & Magnuson, 2009). In contrast, thematic competitors were initially weakly activated and the activation grew steadily late into the course of the trial. This pattern arose because the model was trained to produce thematically-grouped two-word sequences, so even in single word production there was a residual tendency to prepare to produce a thematically related word.

Table 1 shows the normalized summed (over the first 15 time ticks) activation for the taxonomic and thematic competitors for each of the three conditions. The one-way behavioral dissociation is very clear: normalized summed activation for the taxonomic competitors is substantially higher than for the thematic competitors in all three model tests. The model also exhibits the neuroanatomical double dissociation: the event lesion increases the normalized summed activation for the thematic competitors much more than for the taxonomic competitors, and the semantic features lesion increases the normalized summed activation for the taxonomic competitors much more than for the thematic competitors.

A simple model trained to produce two-word sequences based on a stable event representation and a sequential semantic feature-based representation of individual concepts was able to perform single word production. Importantly, the model captured both of the key data patterns from the picture naming VLSM study (Schwartz et al., in press). There was a one-way behavioral dissociation such that




Figure 2. Average activation for target word (circles), taxonomically related words (triangles), and thematically related words (squares) in the trained unlesioned model and two lesion models.

activation of taxonomic responses was higher than activation of thematic responses and a neuroanatomical double dissociation such that damage to semantic feature representations produced increased activation of taxonomic responses relative to thematic responses and damage to event representations produced increased activation of thematic responses relative to taxonomic responses.

Table 1. Normalized summed activations for targets, taxonomically related competitors, and thematically related competitors for each of the three models.

Model	Targets	Taxonomic Competitors	Thematic Competitors
Control	0.703	0.158	0.139
Event Lesion	0.687	0.159	0.154
Sem. Feat. Lesion	0.661	0.181	0.158

Experiment

This experiment was designed to test two behavioral predictions from the simulations reported above: (1) taxonomic and thematic competitors are both activated automatically during single word processing, even when the task demands do not require it, and (2) taxonomic competitors are activated more quickly and more strongly than thematic competitors. These predictions were tested in the domain of spoken word comprehension using the “visual world paradigm” (VWP). In a typical VWP experiment, 4 pictures of objects are shown and the participant is instructed to select a named target object. Previous studies using this paradigm have shown that participants are more likely to look at pictures of objects that are semantically related to the target than at unrelated objects (e.g., Huettig & Altmann, 2005; Mirman & Magnuson, 2009; Yee & Sedivy, 2006), though not at objects that are only related by lexical co-occurrence with no semantic relationship (Yee, Overton, & Thompson-Schill, 2009). The present study specifically distinguished taxonomic semantic similarity and thematic semantic similarity.

Methods

Participants. Fifteen older adult participants (53% females; 27% African American) completed the study. They were selected to be approximately the same age and education level as the aphasic participants in the VLSM study (Schwartz et al., in press). Their mean age was 63 (range: 42-72) and mean years of education was 15 (range: 12-20). All participants had English as their native language and no major psychiatric or neurologic co-morbidities. Mean score on the Mini-Mental State Exam was 29 (range: 26-30). Participants were paid for their participation and reimbursed for travel and related expenses.

Materials. A normed set of 260 color line drawings of common objects (Rossion & Pourtois, 2004) was used for the picture stimuli. Images had a maximum size of 200 x 200 pixels and were scaled such that at least one dimension

was 200 pixels. Critical pairs were selected on the basis of sharing a semantic category (taxonomically related) or frequently participating in an event (thematic relation). Target and competitor words were matched on word frequency, familiarity, length, and neighborhood density across the two conditions (all $p > 0.15$). Stimulus words were recorded by a native English speaker at 44.1kHz. The individual words were edited to eliminate silence at the beginning and end of each sound file.

Apparatus. Participants were seated approximately 24 inches away from a 17-inch monitor with the resolution set to 1024x768 dpi. Stimuli were presented using E-Prime Professional 2.0 experimental design software. Responses were recorded using a mouse. During the testing session, a remote Eyelink 1000 eye tracker was used to record participants’ left eye gaze position at 250 Hz.

Procedure. Each trial was initiated by the participant by clicking on a plus sign (+) in the center of the screen, which caused a four-image display to appear with each image near one of the screen corners. The position of the four pictures was randomized. The display was presented for a 1300ms preview to allow for initial fixations that are driven by visual salience rather than word processing. During the last 300ms of the preview, a red circle appeared in the center of the screen in order to drive the attention back to the neutral central location. After the 1300ms preview, participants heard the target word through speakers and had to click on the image that corresponded to the target word. Each display contained a target object image, a semantic competitor (taxonomic or thematic), and two unrelated distractors. There were a total of 70 trials: 10 practice trials (on which feedback was provided), 20 trials with taxonomic competitors, 20 trials with thematic competitors, and 20 filler trials where none of the images were related to each other.

Results and Discussion

Accuracy was very high ($> 99\%$ correct in both conditions, $p > 0.3$) and mean response times were approximately 2000ms from word onset with no difference between conditions (Taxonomic: $M = 2018$, $SD = 396$; Thematic: $M = 1959$, $SD = 496$; $F < 1$, $p > 0.3$). Only correct response trials were included in the fixation analysis. Figure 3 shows the time course of fixations to the target, semantically related competitor, and unrelated distractors (average of the two distractors) from word onset. Participants were more likely to fixate semantically related competitors than unrelated distractors in both the Taxonomic and Thematic conditions.

To quantify the time course of the semantic competition effects we used Growth Curve Analysis (GCA), a multilevel regression modeling technique using fourth-order orthogonal polynomials (for details see Mirman, Dixon, & Magnuson, 2008). The analysis considered semantic competitor and unrelated distractor fixations from 500ms




Figure 3. The average time course of fixation proportions to the target (T), semantically related competitor (C), and unrelated distractor (U) objects starting at target word onset. Error bars indicate $\pm 1\text{SE}$.

after target word onset (shortly before the target fixations begin to separate from the other conditions, indicating that fixations are starting to be driven by linguistic/semantic processing) to 1700ms after word onset (at which point competition has been mostly resolved and competitor fixations are nearly at floor). The GCA results confirmed semantic competition in both conditions with statistically significant effects of object type (competitor vs. unrelated) on the intercept term (overall more fixations to the semantic competitor than the unrelated distractor) and on the quadratic term (steeper fixation curve rise and fall for the semantic competitor than the unrelated distractor), as well as other, less relevant, model terms (full GCA results are in Table 2).

Table 2. Growth curve analysis results for semantic competition in the two conditions. Parameter estimates are for the semantically related competitor relative to the unrelated distractor.

	Thematic			Taxonomic		
	Est.	<i>t</i>	<i>p</i> <	Est.	<i>t</i>	<i>p</i> <
Intercept	0.036	9.2	0.00001	0.086	12.3	0.00001
Linear	0.037	1.1	n.s.	0.15	2.9	0.01
Quadratic	-0.084	2.3	0.05	-0.24	10.0	0.00001
Cubic	-0.030	2.2	0.05	0.086	6.2	0.00001
Quartic	0.033	2.4	0.05	0.04	3.0	0.01

The semantic competition effect was also substantially larger and peaked earlier in the Taxonomic condition than in the Thematic condition. The Taxonomic competition effect peaked approximately 900ms after target word onset and the Thematic competition effect peaked approximately 1100ms after target word onset. GCA of the full data set examining the interaction of object (competitor vs. unrelated) and condition (Taxonomic vs. Thematic) revealed a clear difference in overall effect size (interaction effect on

intercept term: Est. = -0.049, *t* = 4.88, *p* < 0.0001) and a difference in time course (interaction effect on quadratic term: Est. = 0.156, *t* = 3.17, *p* < 0.01; and on cubic term: Est. = -0.116, *t* = 5.77, *p* < 0.0001; interaction effects on linear and quartic term were not significant). The degree of semantic relatedness contributes to the magnitude and time course of semantic competition effects (e.g., Mirman & Magnuson, 2009), but the current results are also strikingly consistent with the simulations reported above.

In sum, the experiment revealed that thematically and taxonomically related competitors are both activated in the course of spoken word recognition and suggested that taxonomic competitors are activated more quickly and more strongly than thematic competitors. These results are consistent with the predictions of a computational model, which also accounted for behavioral and neuroanatomical results from a large-scale study of aphasic picture naming errors (Schwartz et al., in press).

Conclusions

Based on behavioral, neuroimaging, and VLSM data, Schwartz et al. (in press) proposed that there are two parallel semantic systems. One system, with ATL as a critical hub, captures taxonomic relations based on feature overlap and is particularly important for single object processing and identification. A second system, with TPJ as a critical hub, captures thematic relations based on complementary roles in events and is possibly more relevant for relational processing (e.g., Wu et al., 2007) and sentence processing. We described a computational model that is a concrete implementation of this theory. The model accounted for the one-way behavioral dissociation in aphasic picture naming errors (more taxonomic errors than thematic errors) and the neuroanatomical double dissociation (damaging feature representations leads to relatively more taxonomic errors, damaging event representations leads to relatively more thematic errors).

In addition, the model predicted that both taxonomic and thematic competitors should be automatically activated during single word processing, with taxonomic competitors activated more quickly and more strongly. Results from a spoken word comprehension experiment using eye tracking to assess the time course of competitor activation were consistent with these predictions.

These results suggest that semantic knowledge is represented in two parallel systems – one that is primarily sensitive to semantic feature overlap and taxonomic relations with ATL as the critical hub, and one that is primarily sensitive to event and thematic role relations with TPJ as the critical hub.

Acknowledgements

This research was supported by the Moss Rehabilitation Research Institute and National Institutes of Health grant DC01085 to DM.

References

Binder, J.R., Desai, R. H., Graves, W. W., & Conant, L. L. (2009). Where is the semantic system? A critical review and meta-analysis of 120 functional neuroimaging studies. *Cerebral Cortex, 19*(12), 2767-2796.

Cree, G. S., McRae, K., & McNorgan, C. (1999). An attractor model of lexical conceptual processing: Simulating semantic priming. *Cognitive Science, 23*(3), 371-414.

Estes, Z., Golonka, S., & Jones, L. L. (2011). Thematic thinking: The apprehension and consequences of thematic relations. *Psychology of Learning and Motivation, 54*, 249-294.

Goldwater, M. B., Markman, A. B., & Stilwell, C. H. (2011). The empirical case for role-governed categories. *Cognition, 118*, 359-376.

Hare, M., Jones, M., Thomson, C., Kelly, S., & McRae, K. (2009). Activating event knowledge. *Cognition, 111*(2), 151-167.

Hodges, J. R., Graham, N., & Patterson, K. (1995). Charting the progression in semantic dementia: Implications for the organisation of semantic memory. *Memory, 3*(3-4), 463-495.

Huettig, F., & Altmann, G. T. M. (2005). Word meaning and the control of eye fixation: Semantic competitor effects and the visual world paradigm. *Cognition, 96*(1), B23-B32.

Kalénine, S. N., Peyrin, C., Pichat, C. D., Segebarth, C., Bonthoux, F. O., & Baciu, M. (2009). The sensory-motor specificity of taxonomic and thematic conceptual relations: A behavioral and fMRI study. *NeuroImage, 44*(3), 1152-1162.

Lambon Ralph, M. A., McClelland, J. L., Patterson, K., Galton, C. J., & Hodges, J. R. (2001). No right to speak? The relationship between object naming and semantic impairment: Neuropsychological evidence and a computational model. *Journal of Cognitive Neuroscience, 13*(3), 341-356.

Lin, E.L. & Murphy, G.L. (2001). Thematic relations in adults' concepts. *Journal of Experimental Psychology: General, 130*(1), 3-28.

Mirman, D., Dixon, J. A., & Magnuson, J. S. (2008). Statistical and computational models of the visual world paradigm: Growth curves and individual differences. *Journal of Memory and Language, 59*(4), 475-494.

Mirman, D., & Magnuson, J. S. (2009). Dynamics of activation of semantically similar concepts during spoken word recognition. *Memory & Cognition, 37*(7), 1026-1039.

O'Connor, C. M., Cree, G. S., & McRae, K. (2009). Conceptual hierarchies in a flat attractor network: Dynamics of learning and computations. *Cognitive Science, 33*(1), 1-44.

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.

Pearlmutter, B. A. (1995). Gradient calculation for dynamic recurrent neural networks: A survey. *IEEE Transactions on Neural Networks, 6*, 1212-1228.

Plaut, D. C. (1995). Semantic and associative priming in a distributed attractor network. In J. D. Moore & J. F. Lehman (Eds.), *Proceedings of the 17th Annual Conference of the Cognitive Science Society* (pp. 37-42). Mahwah, NJ: Lawrence Erlbaum Associates.

Rogers, T. T., & McClelland, J. L. (2004). *Semantic cognition: A parallel distributed processing approach*. Cambridge, MA: MIT Press.

Rosson, B., & Pourtois, G. (2004). Revisiting Snodgrass and Vanderwart's object pictorial set: The role of surface detail in basic-level object recognition. *Perception, 33*(2), 217-236.

Schwartz, M.F., Kimberg, D.Y., Walker, G.M., Brecher, A., Faseyitan, O., Dell, G.S., Mirman, D., and Coslett, H.B. (in press). A neuroanatomical dissociation for taxonomic and thematic knowledge in the human brain. *Proceedings of the National Academy of Sciences*.

Schwartz, M. F., Kimberg, D. Y., Walker, G. M., Faseyitan, O., Brecher, A., Dell, G. S., and Coslett, H. B. (2009). Anterior temporal involvement in semantic word retrieval: Voxel-based lesion-symptom mapping evidence from aphasia. *Brain, 132*, 3411-3427.

Wu, D. H., Waller, S., & Chatterjee, A. (2007). The functional neuroanatomy of thematic role and locative relational knowledge. *Journal of Cognitive Neuroscience, 19*(9), 1542-1555.

Yee, E., Overton, E., & Thompson-Schill, S. L. (2009). Looking for meaning: Eye movements are sensitive to overlapping semantic features, not association. *Psychonomic Bulletin & Review, 16*(5), 869-874.

Yee, E., & Sedivy, J. C. (2006). Eye movements to pictures reveal transient semantic activation during spoken word recognition. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 32*(1), 1-14.