Feature-Semantic Gradients in Lexical Categorization Revealed by Graded Manual Responses

Rick Dale (rad28@cornell.edu) **Nicholas C. Hindy (nch24@cornell.edu)**
Michael J. Spivey (spivey@cornell.edu)

Department of Psychology, Cornell University, Ithaca, NY, 14853

Abstract

Participants performed a categorization task in which basic-level animal names (e.g., *cat*) were assigned to their superordinate categories (e.g., *mammal*). Manual motor output was measured by sampling computer-mouse movement while participants clicked on the correct superordinate category label, and not on a simultaneously presented incorrect category. Animal names were selected from the concept-name set of McRae, de Sa, & Seidenberg (1997), in which each concept is associated with a sparse semantic feature vector. If the competing category label draws motor attraction during the categorization task, this attraction should be predicted by feature-semantic measures based on animals' proximity to the incorrect category. This proximity was computed by comparing each animal's feature vector to the mean vector of alternative category choices (e.g., *cat*'s vector to the central tendency of all *reptile* vectors). Dependent measures were computed from mouse-movement trajectories. Degree of trajectory curvature correlated with the proximity of an animal's vector to the mean vector of alternative categories, but only in a particular feature-semantic space. Results suggest that continuous motor output may systematically reflect underlying cognitive processing.

Keywords: Categorization, typicality, semantics, representation, motor output

Introduction

An increasing amount of research reveals that dynamic characteristics of motor output reflect underlying cognitive processing, rather than simply reflecting the discrete decision resulting from that processing. For example, when the cognitive system directs manual output amidst an array of graspable objects, the arm's movement does not always proceed in ballistic fashion toward a single selected object, but may vary continuously depending on the nature of underlying processing. Both manual output and oculomotor responses demonstrate these dynamic characteristics intrinsic to the temporal extent of a response, not just the final outcome of the response. For example, Doyle and Walker (2001) demonstrate that saccadic eye movements reflect attentional processing of visual cues in a simple fixation experiment. Saccade trajectories to the same location exhibit very subtle differential curvature depending on the position of distractor or cue stimuli (see also Sheliga, Riggio, & Rizzolatti, 1995). Additionally, considerable research over the past 10 years has shown that eye movements offer a semi-continuous measure of ongoing cognitive processing (Ballard, Hayhoe, & Pelz, 1995; Tanenhaus, Spivey-Knowlton, Eberhard, & Sedivy, 1995; Underwood, 2005). Aggregate data from eye movements often indicate a graded nature inherent to cognition in general.

Similar findings demonstrate that manual motor output can reveal graded representations. The force and velocity of manual responses vary concomitantly with frequency in a lexical decision task (Abrams & Balota, 1991; Balota & Abrams, 1995), and response and stimulus probability in simple reaction-time tasks (Mattes, Ulrich, & Miller, 2002; Ulrich, Mattes, & Miller, 1999; see also Osman, Kornblum, & Meyer, 1986; Balota, Boland, & Shields, 1989). And in experimental work similar to the saccade trajectory experiments described above, Tipper, Howard, and Jackson (1997) have shown that arm trajectories can curve depending on the visual distractor context in which reaching motions are made (see also Tipper et al., 1992; Sheliga et al., 1997). More recently, Spivey, Grosjean, and Knoblich (2005) and Dale, Kehoe, and Spivey (in press) used computer-mouse trajectories to show that graded manual output reveals temporal continuity in the underlying cognitive processes in spoken word recognition and categorization.

In the latter two studies, manual trajectories were measured through streaming x-y coordinates of computer-mouse movement, and revealed attraction to other response choices in the visual display. For example, in Dale et al. (in press), mouse trajectories were recorded during lexical and pictorial categorization of animal exemplars. Participants categorized an animal by clicking the mouse on one of two category choices. Mouse-movement trajectories consisted of a movement from the bottom center of the screen, to the correct target on the upper left- or right-hand corner of the screen (beside which was a competing category label). Target trials used atypical animals (e.g., *whale*) with an incorrect competitor category that had considerable overlap in terms of semantic and visual features (e.g., *fish*). Though participants responded by clicking the appropriate category (e.g., *mammal*), mouse-movement trajectories exhibited substantial attraction toward the competitor category. Competing activation of the incorrect category in these trials was evident even in the properties of the resultant motor output, and not simply in the decision processes leading up to it. Information flows from the sensors into the categorization process and does not "discretize" before issuing motor output to the effectors. Instead, the effectors themselves seem to reflect some of this processing given the typicality of the exemplar (e.g., *rabbit* vs. *whale*), and the featural overlap with the competing category (e.g., *fish*).

So far, this literature has revealed effects of movement dynamics in simple experimental manipulations. For example, in both the saccade (e.g., Doyle & Walker, 2001) and manual response research (e.g., Tipper et al., 1997), trajectory curvature occurs in simple contexts containing visual distractors. An important outstanding concern is the extent to which properties of motor dynamics reflect finer-grained aspects of the underlying processing task. To make this point clearer, consider the categorization experiments in Dale et al. (in press) just described. While these experiments relied on two groups of animal exemplars, highly typical (e.g., *rabbit*) and highly atypical (e.g., *whale*), they did not explore whether typicality gradients between these extremes are evident in mouse trajectories. Similar research on lexical decision suggests that there should be a relationship between such stimulus parameters and motor output (e.g., Abrams & Balota, 1991). Given the extensive influence of typicality in categorization (see Murphy, 2002), and that motor output may reflect cognitive processing, one should also expect that motor output would reveal gradedness as a function of typicality in a similar categorization task. In the current paper, such gradients are explored through semantic-space measures based on feature norms (McRae et al., 1997).

Moreover, through these feature norms (described below), further details regarding the underlying factors contributing to graded motor output can be acquired by exploring what specific semantic features define the gradients along which output varies. For example, when categorizing animal exemplar names (lexical items), one might expect that certain semantic features constraining that process would exert more of an influence than others, such as visual features if animals were presented using pictures (e.g., visual vs. non-visual properties about the animal).

The following experiment aims to supply some insight into these issues. A large set of animal names is categorized in the same task as Dale et al. (in press), but the competing, incorrect category is randomly selected from 4 possible alternatives. Motor output is again measured in terms of mouse trajectories. The subsequent analysis provides clues about finer-grained processing exhibited by motor output. Firstly, effects akin to typicality gradients should be revealed in the motor output by comparing similarity (or distance) in semantic space between categorized animal exemplars and the competing category. Secondly, because lexical items are being processed, we use multiple feature-semantic measures to reveal that specific feature sets are related to the gradients along which motor output varies. Results demonstrate that the effectors exhibit cognitive processing in systematic ways: Effects found in research on categorization decisions are also revealed in their motor output. In addition, motor output may uncover the semantic features of the stimuli that underlie the lexical categorization task.

Experiment

Participants

31 Cornell University undergraduates participated in the study for extra credit in psychology courses. All participants were right-handed.

Materials

Basic-level animal names were selected from the concept-name set of McRae et al.'s (1997) study in which participants listed features of various animals and objects. For the present study, we used 125 of McRae et al.'s animal names. Each animal corresponded with a superordinate category of mammal, fish, reptile, bird, or insect. The experiment was programmed using RealBasic, and presented on an Apple eMac computer. A standard one-button Apple mouse was sampled using RealBasic's Timer control at a rate of approximately 40 Hz.

Procedure

At the start of each experimental trial, participants were presented with two superordinate animal categories, one category name in the upper right-hand corner of the computer screen and one category name in the upper left-hand corner (with approximately 16 degrees of visual angle between categories). After 2000 ms, a 1cm^2 square appeared at the bottom center of the screen (approximately 13 degrees of visual angle from either category name). When participants clicked on this square with the computer mouse, the square was replaced by a basic-level animal name that corresponded with one of the two super-ordinate animal categories already at the top of the screen. Participants' task was to click on the super-ordinate category corresponding with the animal name for that trial. They were informed to respond naturally and accurately, and were not encouraged to do so in a speeded manner. Before the 125 experimental trials, each participant completed three practice trials. The animal name presentation order and each trial's *incorrect* category were randomized. Likewise, the presentation side of the category names (left vs. right) was also random. Streaming x-y coordinates were recorded between participants' click on the square, and their final categorization choice (see Fig. 1A).

Feature-Semantic Measures

Three different semantic spaces were constructed using sets of features, which composed a unique vector for each animal, and formed the basis for semantic-gradient measures of the proximity between animal and competing category in these spaces. As already mentioned, any motor attraction exhibited in a trial is hypothesized to be due at least partly to the semantic similarity between the incorrect category (serving to subtly attract the manual trajectory) and the animal name that is being categorized.

Each of the 125 animal names can be represented as a sparse semantic vector in a 205-dimension feature space drawn from McRae et al.'s (1997) concept-name set.[1] These semantic features were organized into three groupings: dynamic, static, and category features. Dynamic features (109 total) depicted specific animal behaviors, e.g., "swims" and "eats seeds." Static features (74) depicted specific appearance characteristics, e.g., "has a long tail" and "is furry." Category terms (22) included non-behavioral and non-visual labels often used to classify animals, e.g., "is domestic" and "is endangered." These groupings defined the three following feature-semantic gradients: Proximity in semantic space using dynamic features, static features, and the full 205-feature set.

The 125 animals can be mapped in a semantic space with dimensionality of the number of relevant features (dynamic, static, or full). Each category's central tendency point was determined by averaging the coordinates of its constituent animals. In the resulting space, the most typical animal exemplars (e.g., *rabbit*) of each category clustered around their category's central tendency point. On the other hand, atypical animal exemplars (e.g., *whale*) of each category were positioned much further away, often nearer to the central tendency point of an entirely different animal category (e.g., *fish*).

Manual Measures and Analyses

Streaming x-y coordinates were sampled from the presentation of the lexical item, to the final categorization click, and only correct trials were subjected to analysis. Four properties of manual output were extracted from these mouse-movement data (see Fig. 1B). First, while not used in the main analyses, the movement initiation time was computed by measuring the number of samples before mouse movement was detected (i.e., while the cursor was motionless after the start of a trial). We use this measure of latency below for an additional analysis.

From the remaining trajectory representing output motion, we calculated the number of time steps required to finish the categorization (movement time), the total area occupied by the trajectory compared to an assumed straight line to the correct category (area), and the closest point in the trajectory to the competing category (proximity to incorrect category). The first measure, movement time, was scored using sample counts extracted in RealBasic.[2]

Analyses aimed to detect the extent to which each feature set predicted these dependent measures. In other words, if an animal's feature vector is close to a competing category's vector, one would predict that output measures would represent attraction towards that category label. Therefore, proximity to the incorrect category label should be smaller (closer to label) if the distance between the animal and incorrect category in semantic space is small. To test this, we performed a two-part analysis. Both tests make use of item-based observations. Each animal is paired with 4 randomly selected non-targets. With 125 animals, we therefore have 500 item types supplied by the experiment. The subsequent analyses are based on these 500 item-category pairs. Each pair has observations averaged across participants who encountered it in a trial during the experiment. These data were used in the two-part analysis.

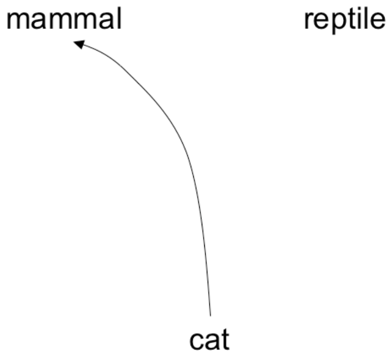

Figure 1A: What participants saw, along with a hypothetical mouse trajectory to the correct category.

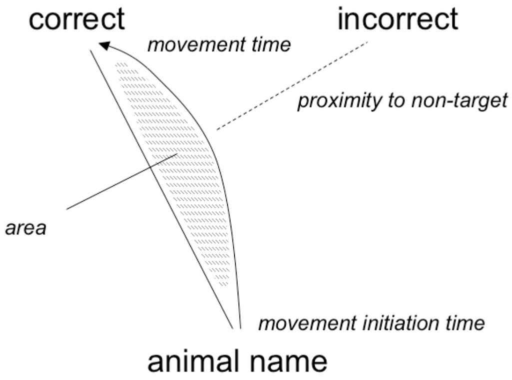

Figure 1B: Depiction of dependent measures.

In the first part, we separated animals along feature-semantic gradients by using one standard deviation (SD) of the mean distance from animals to a given category (i.e., all animals to *bird*). This produced two groups of animals for each category. The first group (N ≅ 60 in the three feature sets), one SD *below* the mean distance, represents those animals *close in semantic space* to the competing category. The second group (N ≅ 60), one SD *above*, is particularly

---

[1] We removed features that occurred uniquely in one animal. Any feature that defined a dimension in the feature-semantic measures was listed in at least 2 or more animal concepts in the McRae et al. (1997) concept set.

[2] The RealBasic Timer function has an approximate sampling rate of 40 Hz or 25 ms. The actual milliseconds at which the mouse is sampled is somewhat different, approximately 33.3ms. For this reason, rather than approximating a derived temporal measure, we represented the results in terms of "ticks" of the RealBasic Timer.

distant from the central tendency of the competing category (see Table 1 for examples). We conducted *t*-tests between these groups for each measure.

In a second test, we sought to confirm that the relationship between feature-semantic distances and the measures holds across the entire group of animals. To do this, we computed a regression coefficient between the semantic measures and the output measures. Proximity to incorrect category, for example, should reveal a positive relationship with semantic distance: Closer semantic distance measures should significantly predict closer spatial proximity measures in the mouse movements.

Table 1: Some examples of close and distant animals in *dynamic* semantic space

Close animals (one SD below mean of Euclidean distance for given category)

animal	category	distance
seal	fish	.30
dove	insect	.31
walrus	fish	.34
Distant animals (one SD above mean of Euclidean distance for given category)		
skunk	insect	1.09
moth	fish	1.08
bull	bird	1.06

Results

Participants erred on 2.45% of experimental trials. These trials were not included in statistical tests.

In comparing the groups of above/below one SD to the mean distance in *dynamic* feature space, animals close to competing categories exhibited larger trajectory area (69386.7 vs. 64671.4 pixels^2, $t(122) = 2.0$, $p < .05$), longer time in motion (26.3 vs. 23.3 samples, $t(122) = 2.6$, $p < .05$), and significantly closer proximity to the competing category label (340.7 vs. 368.1 pixels, $t(122) = 2.6$, $p < .05$). Neither static nor full 205-dimensional feature space exhibited any significant or marginally significant differences.

Regression analyses revealed the same pattern. Only dynamic feature space again revealed significant relationships between output measures and distance. These are presented in Table 2, along with the results for static and full space regressions.

Table 2: Regressions across sets and output measures

	r		
Measure	Dynamic	Static	Full
Area	-.13**	-.01	-.03
In motion	-.15***	.04	.00
Proximity	.13**	.04	.04

*, $p < .05$; ** $p < .01$; *** $p < .001$

Discussion

Results further support that dynamic properties of motor output reflect cognitive processing. The output measures significantly relate to feature-semantic gradients for categories and animals that are not their members: The closer the proximity in semantic space between animal and incorrect category, the greater the attraction of the manual response towards that category label. In addition, by separating the feature space in terms of different semantic content (dynamic vs. static features), we find that the dynamic feature set predicts attraction to the non-target label, while static features do not. Although there are a number of caveats regarding the immediate implications of these semantic spaces (see below), this at least suggests that motor output is reflecting finer-grained featural semantics that underlie lexical categorization in the task.

An additional analysis that may test this claim is to conduct similar tests using Latent Semantic Analysis (LSA; Landauer & Dumais, 1997), a semantic representation scheme based on co-occurrence of lexical items in text. If motor output reflects the processing of lexical items specifically, then we should also observe significant results when computing feature-semantic gradients in terms of LSA measures.

LSA Analysis

LSA measures computed semantic *similarity* (rather than distance or dissimilarity) between animal names and the 4 alternative categories to which they do not belong.[3] We should therefore expect the reverse pattern of results for our dependent measures.

Exactly the same strategy was used to separate two groups of differing distance from mean typicality. These did not produce significant results. However, unlike the regression results above for static, but similarly for the dynamic feature space, LSA significantly predicted all dependent measures: area, $r = .14$, $p < .01$, movement time, $r = .11$, $p < .05$, and proximity, $r = -.10$, $p < .05$. While the SD separation of animal-category pairs did not attain significance, the regression results reveal that gradient effects hold with LSA measures.

Movement Initiation Time

We present a final analysis that tests a prediction made by the perspective that processing flows into the effectors. If cognition indeed does not discretize information prior to initiating motor output during categorization, then there is likely an important temporal component to the process. If a participant allows a relatively large amount of time to pass during a trial before initiating her response, then one would not expect there to be significant dynamical competition in the output: By spending more time evaluating the animal name and category labels before moving, the decision

---

[3] We used online LSA tools located at lsa.colorado.edu. We used the text "General reading to 1st year of college" with 300 factors, though most large English texts work for this analysis.

process may reach a higher level of certainty. The upshot may be a more reliably linear, ballistic movement to the correct category label.

We looked at the relationship between movement initiation time and the output measures. One measure exhibits a significant relationship. Proximity to competitor is positively related ($r = .15$, $p < .001$). In other words, the longer the amount of time spent before initiating motor movement, the less spatial attraction exerted by the competing category.

General Discussion

In everyday life, our arms move continuously during such tasks as gesturing in conversation, organizing objects on a table, and managing cooking ingredients. Their neural substrate is a fairly slow system (relying heavily on prediction; e.g., Flanagan & Lolley, 2001), not firing off movements in staccato fashion (as with saccades), but often changing course mid-path, or issuing graded movements as it directs the arms to their target. This intuition about everyday movement is demonstrated in the foregoing results, and in the array of motor-dynamics findings reviewed above. Even in a relatively "higher-order" cognitive process such as categorization, manual output has internal characteristics that likely reflect the categorization process itself. In the above results, mouse trajectories vary concomitantly with semantic gradients, and these gradients may lie along dimensions relevant to the processing task.

Nevertheless, a number of important limitations should be noted. First, the results, while robust, are thus far fairly weak. There may be a number of reasons for this. Previous findings with saccadic trajectories show an effect of location of distractors relative to targets, resulting in varying strengths of trajectory curvature (see Godijn & Theeuwes, 2002, for a review). In the kinds of experiments reported here, it is uncertain where or whether there are effects of relative location. Further studies may explore different locations of competing category labels, and whether this weak result is inherent to the nature of interaction between cognition and action, or perhaps the design presented here involved response choices situated too close or too far to reveal more marked trajectory effects.

Second, little was done to transform the semantic feature space afforded by McRae et al.'s (1997) concept set (e.g., multidimensional scaling, row/column normalization, similarity-metric transformation). Also, these feature sets were not intended to define categories – so raw feature values in Euclidean space were used to infer category clusters. Moreover, we did not make use of semantic scores for correct categories. We feel that this is in fact a more conservative test of the predictions made above, because raw Euclidean distance between animal and incorrect category relates to motor measures, without adding the additional information regarding proximity to correct category semantics. Further detailed analysis of McRae et al.'s (1997) semantic feature space may thus strengthen these results (e.g., using the coding criteria employed by Cree and McRae, 2003).

Despite these limitations, the proximity of incorrect category labels did produce dynamic motor movement effects that reliably correlate with raw semantic feature space. The results further contribute to a wide literature on processing distinctions between mode of stimulus presentation: Categorization of lexical items may rely on semantic information that is distinct from that centrally involved in categorizing pictures of animal exemplars (e.g., Snodgrass, 1984; Viswanathan & Childers, 2003). Although it is possible that the limitations may have rendered detection of static semantic features undetectable, further exploration may seek to explore the contribution of static visual (or other perceptual) information in both decision- and output-based measures of lexical categorization (see, e.g., Pulvermüller, 1999). One approach is to make use of picture stimuli of the 125 animal names (e.g., Dale et al., in press). Another, as mentioned, is to subject the semantic feature space used here to more detailed analyses, perhaps revealing the relevance of detailed visual or other perceptual semantic features in more sensitive tests.

The findings reported here challenge the common intuition that the properties of motor output are uninformative of cognition. Perhaps more importantly, they suggest that processing flows in systematic ways into motor behaviors, rather than simply being collapsed onto them to generate a categorical response (cf. Gold & Shadlen, 2000). They may even recommend a "cascadic flow" perspective on cognition that sees information flow continuously from sensors to effectors (McClelland, 1979; Balota & Abrams, 1995; Spivey et al., 2005).

References

Abrams, R., & Balota, D. (1991). Mental chronometry: Beyond reaction time. *Psychological Science*, *2*, 153-157.

Allopenna, P., Magnuson, J., & Tanenhaus, M. (1998). Tracking the time course of spoken word recognition using eye movements: Evidence for continuous mapping models. *Journal of Memory and Language*, *38*, 419-439.

Balota, D., & Abrams, R. (1995). Mental chronometry: Beyond onset latencies in the lexical decision task. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*, 1289-1302.

Balota, D., Boland, J., & Shields, L. (1989). Priming in pronunciation: Beyond pattern recognition and onset latency. *Journal of Memory and Language*, *28*, 14-36.

Cree, G., & McRae, K. (2003). Analyzing the factors underlying the structure and computation of the meaning of *chipmunk, cherry, chisel, cheese, and cello*, (and many other such concrete nouns). *Journal of Experimental Psychology: General*, *132*, 163-201.

Dale, R., Kehoe, C., & Spivey, M. (in press). Graded motor responses in the time course of categorizing atypical exemplars. *Memory and Cognition*.

Doyle, M., & Walker, R. (2001). Curved saccade trajectories: Voluntary and reflexive saccades curve away from irrelevant distractors. *Experimental Brain Research*, *139*, 333-344.

Flanagan, J.R. & Lolley, S. (2001). The inertial anisotropy of the arm is accurately predicted during movement planning. *Journal of Neuroscience*, *21*, 1361-1369.

Gaveau, V., Martin, O., Prablanc, C., Pelisson, D., Urquizar, C., & Desmurget, M. (2003). On-line modification of saccadic eye movements by retinal signals. *Neuroreport*, *14*, 875-878.

Godijin, R., & Theeuwes, J. (2002). Programming of endogenous and exogenous saccades: Evidence for a competitive integration model. *Journal of Experimental Psychology: Human Perceptiona and Performance*, *28*, 1039-1054.

Gold, J., & Shadlen, M. (2000). Representation of a perceptual decision in developing oculomotor commands. *Nature*, *404*, 390-394.

Landauer, T., & Dumais, S. (1997). A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, *104*, 211-240.

Mattes, S., Ulrich, R., & Miller, J. (2002). Response force in RT tasks: Isolating effects of stimulus probability and response probability. *Visual Cognition*, *9*, 477-501.

McClelland, J. (1979). On the time relations of mental processes: An examination of systems of processes in cascade. *Psychological Review*, *86*, 287-330.

McMurray, B., Tanenhaus, M., Aslin, R., & Spivey, M. (2003). Probabilistic constraint satisfaction at the lexical/phonetic interface: Evidence for gradient effects of within-category VOT on lexical access. *Journal of Psycholinguistic Research*, *32*, 77-97.

McRae, K., de Sa, V., & Seidenberg, M. (1997). On the nature and scope of featural representations of word meaning. *Journal of Experimental Psychology: General*, *126*, 99-130.

Murphy, G. (2002). *The big book of concepts*. Cambridge, MA: The MIT Press.

Osman, A., Kornblum, S., & Meyer, D. (1986). The point of no return in choice reaction time: Controlled and ballistic stages of response preparation. *Journal of Experimental Psychology: Human Perception and Performance*, *12*, 243-258.

Pulvermueller, F. (1999). Word's in the brain's language. *Behavioral and Brain Sciences, 22*, 253-336.

Sheliga, B., Craighero, L., Riggio, L. & Rizzolatti, G. (1997). Effects of spatial attention on directional manual and ocular responses. *Experimental Brain Research*, *114*, 339-351

Sheliga, B., Riggio, L., & Rizzolatti, G. (1995). Spatial attention and eye movements. *Experimental Brain Research*, *105*, 261-275.

Snodgrass, J.G. (1984). Concepts and their surface representations. *Journal of Verbal Learning and Verbal Behavior*, *23*, 3-22.

Spivey, M., Grosjean, M., & Knoblich, G. (2005). Continuous attraction toward phonological competitors. *Proceedings of the National Academy of Sciences*, *102*, 10393-10398.

Spivey, M., Tanenhaus, M., Eberhard, K., & Sedivy, J. (2002). Eye movements and spoken language comprehension: Effects of visual context on syntactic ambiguity resolution. *Cognitive Psychology*, *45*, 447-481.

Tipper, S., Howard, L., & Jackson, S. (1997). Selective reaching to grasp: Evidence for distractor interference effects. *Visual Cognition*, *4*, 1-38.

Ulrich, R., Mattes, S., & Miller, J. (1999). Donders's assumption of pure insertion: An evaluation on the basis of response dynamics. *Acta Psychologica*, *102*, 43-75.

Underwood, G. (2005) (Ed.), *Eye Guidance in Cognition*. Oxford: Oxford University Press.

Viswanathan, M. & Childers, T.L. (2003). An enquiry into the processing of categorization of pictures and words. *Perceptual and Motor Skills*, *96*, 267-287.