

COMMENTARY

Clarifying the Nature of the Distinctiveness by Domain Interaction in Conceptual Structure: Comment on Cree, McNorgan, and McRae (2006)

Kirsten I. Taylor, Angeliki Salamoura, Billi Randall, Helen Moss, and Lorraine K. Tyler
University of Cambridge

The conceptual structure account of semantic memory (CSA; L. K. Tyler & H. E. Moss, 2001) claims that feature correlation (the degree to which features co-occur) and feature distinctiveness (the number of concepts in which a feature occurs) interact with domains of knowledge (e.g., living vs. nonliving) such that the distinctive features of nonliving things are more highly correlated than the distinctive features of living things. Evidence for (B. Randall, H. E. Moss, J. M. Rodd, M. Greer, & L. K. Tyler, 2004) and against this claim (G. S. Cree, C. McNorgan, & K. McRae, 2006) has been reported. This comment outlines the CSA, discusses Cree et al.'s (2006) critiques of the Randall et al. (2004) experiments and the CSA, and reports new analyses of property norm and behavioral data, which replicate the results reported by Randall et al. (2004).

Keywords: conceptual structure, distinctiveness, correlation, domain, semantic memory

Brain-damaged patients with category-specific semantic impairments in one category of knowledge provide invaluable clues about how conceptual knowledge is organized in the brain. The most common variant of this disorder is a semantic impairment that disproportionately impacts concepts belonging to the domain of living things, with a relative sparing of semantic information about nonliving concepts (Humphreys & Forde, 2001). One set of models developed to explain this pattern of semantic deficits proposes that neuroanatomically distinct regions store different categories of knowledge (e.g., the categories of living animate, living inanimate, conspecifics, and perhaps tools; Caramazza & Mahon, 2003, 2005; Caramazza & Shelton, 1998) or the features that are relevant in the representation of living and nonliving things. For example, visual semantic features are claimed to be central to the representation of living things, whereas functional semantic features are important to the representation of nonliving things (Warrington & McCarthy, 1983, 1987; Warrington & Shallice, 1984; see also Martin & Chao, 2001, and Martin, Ungerleider, & Haxby, 2000, for a more recent conceptualization).

Another class of models is based on distributed accounts of conceptual representations (e.g., Masson, 1995; McRae, de Sa, & Seidenberg, 1997). These models claim that concepts are represented in a distributed system composed of units corresponding to semantic features and that differences in the distributional properties of features comprising living and nonliving concepts make them differentially susceptible to the effects of brain damage (Devlin, Gonnerman, Andersen, & Seidenberg, 1998; Durrant-Peatfield, Tyler, Moss, & Levy, 1997; Gonnerman, Andersen, Devlin, Kempler, & Seidenberg, 1997; Tyler, Moss, Durrant-Peatfield, & Levy, 2000). One instantiation of this type of model is the conceptual structure account (CSA; Greer et al., 2001; Moss, Tyler, & Devlin, 2002; Tyler & Moss, 2001; Tyler et al., 2000). A key claim of the CSA is that feature correlation (the degree to which features co-occur) and feature distinctiveness interact with living–nonliving domains, such that distinctive features of nonliving things are more highly correlated than the distinctive features of living features. Because highly correlated features are thought to be spared by brain damage, and highly distinctive features are necessary for the identification of a concept and its differentiation from other similar concepts, this interaction explains the relatively greater prevalence of category-specific semantic deficits for living things whose weakly correlated distinctive features are more susceptible to brain damage.

In a recent article, we reported findings from a property norm study with healthy individuals that supported the CSA's predicted interaction of distinctiveness and correlation with domain (Randall, Moss, Rodd, Greer, & Tyler, 2004). In a speeded feature verification task, healthy participants processed distinctive and correlated nonliving features more quickly than distinctive, less correlated living features (Randall et al., 2004). This study used the most distinctive features, that is, those occurring in one or two concepts in a set of previously collected property norms.

Kirsten I. Taylor, Angeliki Salamoura, Billi Randall, Helen Moss, and Lorraine K. Tyler, Centre for Speech, Language and the Brain, Department of Experimental Psychology, University of Cambridge.

The writing of this article was supported by grants from the Medical Research Council and the British Academy to Lorraine K. Tyler and a Marie Curie Intra-European Fellowship to Kirsten I. Taylor.

We thank George Cree and Ken McRae for their reviews of the manuscript, Ken McRae for providing us with the data necessary to calculate the correlational strengths of the features in the norms obtained by his lab, and Mike Aitken for his statistical advice.

Correspondence concerning this article should be addressed to Lorraine K. Tyler, Department of Experimental Psychology, University of Cambridge, Downing Street, Cambridge, England, United Kingdom CB2 3EB. E-mail: lktyler@csl.psychol.cam.ac.uk

The basis of these results was questioned in a recent article by Cree et al. (2006). In contrast to Randall et al. (2004), Cree et al. found that the most distinctive features of living and nonliving things were processed equally quickly and that both were processed more quickly than shared features. These results were interpreted as supporting Cree et al.'s feature-based connectionist attractor network model and disconfirming the CSA. However, Cree et al.'s experiments did not analyze correlations of most distinctive features, as they argued that correlations involving such distinctive properties are spurious and not psychologically meaningful. A major claim of the CSA is that both distinctiveness and correlation determine conceptual structure. Thus, the more distinctive features of living and nonliving things are represented and processed differently by virtue of their correlational status, not their domain membership per se. Because Cree et al. did not calculate the correlational strengths of their distinctive living and nonliving features, these stimuli cannot be used to test the CSA's claim. In this comment, we outline the CSA and discuss the various critiques that Cree et al. made of the Randall et al. article in particular and the CSA in general. We also report new analyses of the Randall et al. property norm and behavioral data and show that the pattern of results remains unchanged. Finally, we discuss Cree et al.'s behavioral study in the context of the CSA and provide an alternative view of their findings.

The Conceptual Structure Account

The CSA was developed to explain the relatively greater prevalence of category-specific semantic impairments for living compared with nonliving things reported in the neuropsychological literature and the performance patterns of patients with semantic impairments for living things (Moss et al., 2002; Moss, Tyler, Durrant-Peatfield, & Bunn, 1998; Tyler & Moss, 2001). The model assumes a distributed, feature-based system in which a concept is represented by its constituent features. The CSA posits that at least two statistical properties of features organize conceptual space, and both determine how concepts are organized in the normal system and which aspects of concepts are lost as a consequence of brain damage. The first is *feature correlation*, that is, the extent to which two features co-occur (Cree & McRae, 2003; Devlin et al., 1998; McRae, Cree, Westmacott, & de Sa, 1999; McRae et al., 1997; Rosch, 1978; Tyler et al., 2000). For example, the feature "has eyes" commonly co-occurs with other features ("has ears," "can see"), the entire set being co-activated upon each encounter with an object. Within a Hebbian-like framework, this mutual co-activation is thought to strengthen the connections between the features, making strongly correlated features more resilient to the effects of brain damage (Gonnerman et al., 1997). Conversely, features that do not typically co-occur (e.g., "has a hump," "has a nose") do not benefit from a strengthened association and are more vulnerable to the effects of brain damage. The second dimension is *distinctiveness*, which is a measure of the number of concepts in which a feature occurs. Whereas some features are shared by many concepts (e.g., "has eyes"), others are not and thus are relatively more distinctive (e.g., "has stripes"). Distinctive features are important, as they are typically more informative about the identity of a concept and thus relatively more important for identifying an object and differentiating it from other similar objects.

The CSA claims that natural differences in the statistical properties of concepts in the living and nonliving domains can explain the relatively higher frequency of category-specific semantic deficits for living compared with nonliving things and the performance patterns of patients with this type of deficit. Specifically, we have claimed, primarily on the basis of property norm data, that living things can be characterized by large clusters of mainly shared features, which are also highly correlated. In contrast, the distinctive features of living things are only weakly correlated with the concept's other features, making these susceptible to the effects of brain damage. Most nonliving things, on the other hand, have smaller clusters of features with relatively more distinctive features, which are also more highly correlated with one another. Randall et al. (2004) summarized these relationships as follows: "The [CSA] claims that the distinctive features of living things (e.g., *a lion's mane*) are only weakly correlated with other features, whereas the shared features are strongly correlated. For nonliving things, both shared and distinctive features tend to be strongly correlated, although they are relatively fewer in number" (p. 394). Thus, the central prediction of the CSA is that the distinctive features of living things, which are less correlated, will be more susceptible to brain damage than the distinctive features of nonliving things, which are more highly correlated. Because the shared properties of both living and nonliving things tend to be spared in patients with category-specific semantic deficits, this supports the CSA's claim that living and nonliving shared features are both strongly correlated. As distinctive features are relatively more important for concept identification, the selective loss of the distinctive features of living things will pattern as a category-specific semantic impairment for living things (Tyler & Moss, 2001).

We have reported neuropsychological evidence of a relative disadvantage for the distinctive features of living things compared with all other feature types in patients with category-specific semantic deficits for living things, consistent with the CSA's primary prediction (Moss et al., 1998; Tyler & Moss, 2001; see Taylor, Moss, & Tyler, 2007, for an overview). A particularly compelling example of this selective deficit is the drawing performance of a patient with herpes simplex encephalitis and a category-specific semantic deficit for living things (Moss, Tyler, & Jennings, 1997). This patient's drawings of artifacts contained both the shared and distinctive features necessary to identify them (see Figure 1a). In striking contrast, the patient's drawings of animals contained only the shared properties characteristic of this category (e.g., four legs, tails, eyes) and were devoid of the distinctive features required to uniquely identify each animal (see Figure 1b).

To determine whether the CSA's main predictions are also supported in the normal conceptual system, we performed a property norm study and a speeded feature verification study with healthy participants (Randall et al., 2004). In the property norm study, we included concepts from the living (31 animals, 16 fruits) and nonliving (22 tools, 24 vehicles) domains (as well as 47 filler items; Randall et al., 2004, see also Greer et al., 2001) and asked participants to list as many different features as possible that belonged to each concept. We used these data to calculate the statistical properties of features belonging to living and nonliving concepts, in particular their correlation and distinctiveness values. We predicted that nonliving things would have relatively more

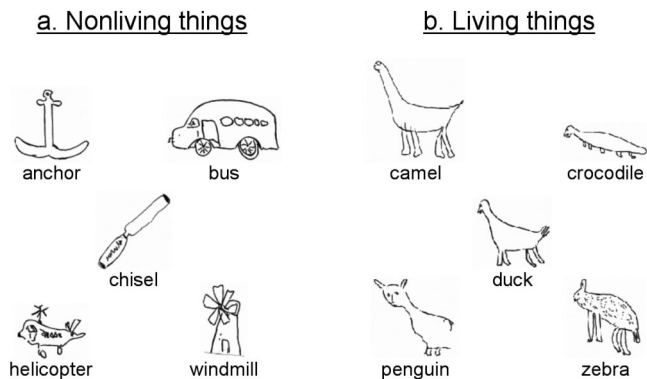


Figure 1. Drawings of (a) nonliving and (b) living things from a patient with herpes simplex encephalitis and a category-specific semantic deficit for living things. From “When Leopards Lose Their Spots: Knowledge of Visual Properties in Category-Specific Deficits for Living Things,” by H. E. Moss, L. K. Tyler, and F. Jennings, 1997, *Cognitive Neuropsychology*, 14, pp. 935, 937. Copyright 1997 by Taylor & Francis. Reprinted with permission.

distinctive features (i.e., properties occurring in one or two concepts) than living things, whereas living things would have relatively more correlated properties. Critically, we predicted that the distinctive features of nonliving things would be more highly correlated than the distinctive features of living things, providing the basis for the relative sparing of the distinctive features of nonliving compared with living things following brain damage. The differences in the distinctiveness and correlation values of living and nonliving things supported these predictions: The mean correlational strengths of the distinctive features of nonliving things was significantly greater than those for living things (.58 vs. .50, respectively), $t(369) = 6.70, p < .001$; Randall et al., 2004).

In a second experiment, Randall et al. (2004) used a speeded feature verification task, with concept–feature pairs derived from the property norms, to determine whether distinctiveness and correlation affect conceptual processing. Assuming that increasing feature correlation increases the speed with which features are initially activated in on-line comprehension tasks (McRae et al., 1997, Randall et al. (2004) predicted that the distinctive features of living things, which are less strongly correlated, would be at a processing disadvantage compared with the distinctive features of nonliving things and the shared features of both living and nonliving things. Shared and distinctive features were selected from the property norm data.¹ Table 1 shows the correlational measures for these features, which establish the greater correlational strengths for distinctive features belonging to nonliving compared with living concepts.

Randall et al. (2004) used a speeded feature verification task in which participants were encouraged to respond rapidly. This was considered to be a crucial aspect of the experiment, as its aim was to tap into the early rise time in feature activation, which McRae et al. (1997) had shown to be influenced by feature correlation. Thus, participants were instructed to “beat the beep” by responding as quickly as possible to concept–feature pairs, where concept words were followed by distinctive or shared features, then a beep. Randall et al. (2004) collected data on the psycholinguistic properties of the selected concepts and features: word length, lemma

frequency, familiarity, imageability, number of senses, production frequency of the feature from the feature production norms, and rated semantic relatedness and association strength of the concept–feature pairs. As stimuli could not be matched across domain and distinctiveness conditions on all of these variables without constructing highly unusual sets, Randall et al. performed additional statistical analyses with all of the psycholinguistic variables influencing performance as covariates (i.e., using analyses of covariance and multiple regression analyses). Bivariate correlations showed that production frequency, rated semantic relatedness, and the association strength between the concept and feature words significantly correlated with feature verification reaction times (RTs), and these variables were therefore entered as covariates in an analysis of covariance. This analysis showed no main effect of domain, but revealed the critical interaction of distinctiveness with domain, reflecting slowest RTs to the distinctive features of living things, an interaction that also accounted for the significant main effect of distinctiveness (with faster RTs to shared than distinctive features). This pattern was partly confirmed with an analysis of covariance of the error rates (with production frequency and rated semantic relatedness as covariates): There was no significant main effect of domain, but a main effect of distinctiveness, with more errors to distinctive than to shared features. Whereas the interaction of domain with distinctiveness did not reach significance, more errors were made with the distinctive compared with shared features of living things. Finally, multiple regression analyses of production frequency, rated semantic relatedness, association strength, and distinctiveness values on feature verification RTs confirmed a significant predictive effect of distinctiveness overall, which, when analyzed separately by domain, was significant for the living but not the nonliving concept–feature pairs. These results support the prediction that the distinctive features of living things are disadvantaged compared with their shared features, as well as the distinctive and shared features of nonliving things, which we claim arises from their impoverished correlation with other features. We now turn to Cree et al.’s (2006) specific critique of the Randall et al. article.

Cree et al. (2006): Critique of Randall et al. (2004)

Correlation and Distinctiveness in Conceptual Structure

Cree et al. (2006) argued that correlations involving distinctive features are spurious and may not meaningfully influence the speed of conceptual processing. This is because distinctive features occur too infrequently to provide reliable estimates of feature co-occurrence. For example, a feature that occurs only once could co-occur with another feature that occurs only once, producing a perfect correlation of 1.0, or features that occur only once or twice can be significantly highly correlated with another feature even if this feature occurs in more than a dozen other concepts, in which case correlation values would range from 1.00 to .21 (i.e., 1.00 if features A and B occur only once and in the same concept, and .21 if feature A occurs only once in concept X and feature B occurs in concept X plus 18 other concepts). The issue of whether the correlational status of distinctive features is important in concep-

¹ Because of an error, the correlational values were misreported in Randall et al. (2004); see Table 1 for correct values.

Table 1
Correlational Strength Values for the Stimuli Used in Randall et al.'s (2004) Behavioral Experiment Data

Stimuli	N	Correlational strength
		M (SD)
Living		
Distinctive	13	.42 (.05)
Shared	30	.35 (.12)
Nonliving		
Distinctive	17	.54 (.13)
Shared	19	.30 (.13)

tual organization and processing is a question that has not yet been experimentally examined. Nevertheless, correlations associated with distinctive features may reflect a different psychological phenomenon than correlations of clusters of shared features, and for this reason we reanalyzed the Randall et al. (2004) property norm data, excluding all distinctive properties to determine whether the most distinctive nonliving features were still more correlated than the most distinctive living features.

Following Cree et al. (2006), we excluded distinctive features (those occurring in only one or two concepts) and calculated the correlation values of the remaining most distinctive features—those occurring in 3 or 4 concepts—in the living and nonliving domains. For this analysis, we used a measure of correlation—intercorrelation—used by Cree et al. (Cree & McRae, 2003). Here, intercorrelation values for the most distinctive features of living things were greater than the intercorrelation values of distinctive nonliving features, $t(117) = 3.36, p = .001$, contrary to the CSA's prediction (see Table 2).

In our research, we have used a different measure of correlation—correlational strength—based on the mean correlation of all significantly correlated property pairs. This measure contrasts with the measure of intercorrelation used by Cree and McRae (2003), which is based on the sum of the variance accounted for by each significantly correlated property pair. Whereas both measures estimate the strength with which a given feature co-occurs with other features, the intercorrelation measure additionally reflects the number of correlated property pairs within a concept: The greater the number of correlated property pairs, the larger is the intercorrelation measure, all other things being equal. What complicates the picture is that living things have larger numbers of features than nonliving things (mean number of features in living and nonliving concepts in the Randall et al., 2004, property norms: $M = 14$ vs. 9, respectively; $t(91) = 7.52, p < .0001$; see also McRae & Cree, 2002, and Devlin et al., 1998). This difference in number of features appears related to the well-established finding that living things have more correlated property pairs than nonliving things (Keil, 1987; Malt & Smith, 1984; McRae et al., 1997; Randall et al., 2004). The more features a concept has, the more likely it is that these are correlated, all other things being equal. Thus, intercorrelation values may be greater for the properties of living things than for those of nonliving things (including more distinctive features) because living things have more features and more correlated features than nonliving things, not because these features are more strongly correlated (see Table 2). Because the

semantic effects of the number of features (see, e.g., Pexman, Holyk, & Monfils, 2003, and Pexman, Lupker, & Hino, 2002) or the number of correlated property pairs associated with a concept may be independent of the effect associated with the strength with which these features are correlated, we prefer to use the correlational strength measure.

The correlational strengths for distinctive living and nonliving features (i.e., those occurring in three or four concepts) and those that are shared are listed in Table 2. Using this measure, we find that the correlational strengths of distinctive nonliving features were more strongly correlated than the distinctive living features, as predicted by the CSA, $t(117) = 2.67, p < .01$. Therefore, Randall et al. (2004)'s report that the correlational strengths of the most distinctive properties of nonliving things was greater than the most distinctive properties of living things does not appear to be due to potentially spurious correlations associated with highly distinctive features.

Psycholinguistic Variables

Cree et al. (2006) also made a number of methodological critiques of the behavioral experiment in Randall et al. (2004), which focused on the matching of psycholinguistic variables. Cree et al. claimed, in particular, that we had not matched on production frequency, association strength, and feature length and frequency and that this could have accounted for or influenced the pattern of results. We are baffled by these comments, as we did, indeed, consider these variables in the analyses. Although it was not possible to adequately match these variables across the sets (because of inherent differences), we included the relevant analyses in which we covaried out the influence of any psycholinguistic variable affecting performance in our covariance and multiple regression analyses. Cree et al. also claimed that the feature verification error rate for the distinctive features of living things was "extremely high." (p. 645). Indeed, the main reasons for using the speeded feature verification task was not only to tap early rise times in feature activation but also as a challenging task that generated enough errors so that they could be meaningfully statistically analyzed. As such, the "extremely high" error rate for the

Table 2
Intercorrelation and Correlational Strength Values for Distinctive Features (Occurring in Three or Four Concepts) and Shared Features (Occurring in More Than Five Concepts) in the Living and Nonliving Domains

Stimuli	N	Intercorrelation	Correlational strength
		M (SD)	M (SD)
Living			
Distinctive	74	104 (71)	.38 (.09)
Shared	359	131 (85)	.43 (.11)
Nonliving			
Distinctive	45	65 (43)	.44 (.16)
Shared	139	80 (54)	.45 (.12)

Note. Analyses exclude distinctive features occurring in one or two concepts.

distinctive features of living things supports our prediction that these features are selectively disadvantaged.

Property Norm Set

Cree et al. (2006) further argued that the small number of concepts in Randall et al.'s (2004) property norm study (93 target concepts and 47 filler items) overestimated a feature's distinctiveness. Cree et al. viewed this as a major weakness of the Randall et al. (2004) study: "[Ninety-three concepts] is a small sample, and it is possible therefore that some features that were distinctive in their norms may actually occur in many more concepts in the world. Therefore, these features may not have been distinctive to participants" (p. 645). We believe that to test the CSA's predictions, we only need relative values on the distinctiveness dimension; the absolute distinctiveness values within the entire conceptual space are not necessary (nor would they be possible to obtain!). Thus, the demonstration of relative differences in the conceptual structure of living and nonliving things, even in a small sample of concepts, is sufficient to support our predictions, provided there is no sampling bias in the concepts selected for the norming study.

To evaluate this hypothesis, we compared correlational strength values obtained from the Randall et al. (2004) property norms with corresponding correlational strength values in McRae, Cree, Seidenberg, and McNorgan's (2005) much larger set of 541 concepts. These comparisons were performed on the 360 concept–feature pairs that were identical in both property norms. As expected, distinctiveness values (defined as 1 divided by the number of concepts in which a feature occurs) were significantly higher in the Randall et al. than in the McRae et al. (2005) property norms ($M = .132$, $SD = .096$ vs. $M = .087$, $SD = .130$, respectively), $F(1, 359) = 61.66$, $p < .001$, as were feature correlation values (feature correlation strength: Randall et al., 2004: $M = .43$, $SD = .11$; McRae et al., 2005: $M = .24$, $SD = .09$, $F(1, 359) = 790.33$, $p < .001$; feature intercorrelation: Randall et al.: $M = 109$, $SD = 76$; McRae et al.: $M = 53$ ($SD = 61$), $F(1, 359) = 174.16$, $p < .001$). Importantly, both distinctiveness values and correlational strengths in the Randall et al. and McRae et al. norms were significantly correlated: distinctiveness, $r(360) = .703$, $p < .001$; feature correlation strength, $r(360) = .296$, $p < .001$; and feature intercorrelation, $r(360) = .420$, $p < .001$.² Thus, it does not appear that the Randall et al. property norms suffered from significant biases from the sampling of concepts or a relatively smaller number of concepts, or, alternatively, both the McRae et al. (2005) and Randall et al. (2004) property norms may suffer from sample bias.

Correlation and Speed of Processing

Cree et al. (2006) argued that the correlation values of the most distinctive features are spurious and may not meaningfully influence speed of processing. We tested for this possibility in Randall et al.'s (2004) behavioral results by recalculating correlational strengths of all significantly correlated feature pairs ($p < .05$) on the basis of production frequency feature vectors, which excluded all unique features (i.e., those occurring in only one concept). Only unique features were excluded from this analysis, as the group of distinctive features used in Randall et al.'s behavioral study included features that occurred in only one or two concepts, and the

Table 3

Correlation Strength and Distinctiveness Values for Living and Nonliving Stimuli Used in Randall et al.'s (2004) Behavioral Experiment Excluding Unique Features and Corresponding Retransformed Inverse Transformed Reaction Times (RTs) and Error Proportions

Stimuli	N	Correlational strength	Distinctiveness	RTs	Error proportion
		M (SD)	M (SD)	M (SD)	M (SD)
Living					
Distinctive	4	.31 (.06)	.50 (.00)	553 (65)	.33 (.20)
Shared	30	.35 (.13)	.16 (.10)	505 (54)	.22 (.14)
Nonliving					
Distinctive	9	.40 (.13)	.50 (.00)	475 (41)	.10 (.12)
Shared	19	.28 (.20)	.15 (.09)	506 (57)	.20 (.15)

exclusion of all distinctive features would have left no stimuli in the group. The mean correlational strength values still show the critical distinction, with distinctive features of nonliving things being more strongly correlated than the distinctive features of living things (see Table 3).

We also recalculated the mean reaction time (RT) and error proportion for each condition in Randall et al.'s (2004) experiment, excluding responses to unique features. Again, the patterns of RTs and error proportions across the four conditions remained the same, with slowest responses and most errors to the distinctive features of living things and fastest responses and fewest errors to the distinctive features of nonliving things. Thus, behavioral performance in the speeded feature verification task still mirrors the pattern of correlational strengths across the more distinctive features of living and nonliving things, even when unique features are excluded from the analyses.

After excluding unique features, only four living features and nine nonliving features remained in the distinctive group. Given these small sets, any statistical analyses can only be suggestive. Nevertheless, we conducted two separate 2×2 ANCOVAs, one with inverse transformed RTs and one with error proportions, with the variables domain (living vs. nonliving) and distinctiveness category (distinctive vs. shared) and the covariates production frequency, rated semantic relatedness, and the association strength between the concept and feature words (the latter for inverse transformed RTs only). These ANCOVAs tested the CSA's prediction that the interaction of domain with correlation and distinctiveness significantly influence the speed of conceptual processing.

Despite the small numbers of stimuli in each condition, these ANCOVAs showed trends consistent with our original results. First, the critical Domain \times Distinctiveness interaction showed a trend in the predicted direction, $F(1, 55) = 3.22$, $p = .078$, and $F(1, 56) = 3.04$, $p = .087$, for inverse transformed RTs and error proportions, respectively. These were driven by slower RTs and a greater proportion of errors for the distinctive features of living things compared with the distinctive features of nonliving things (Mann–Whitney $U = 6.00$, $p = .064$, and Mann–Whitney $U = 4.00$, $p < .05$, respectively; see Table 3). The main effects of

² We thank Ken McRae for providing us with the feature production vectors required to perform these comparisons.

domain and distinctiveness were not significant (all $F_s < 2.3$). Because the distinctive features of living and nonliving things in this data set were perfectly matched on distinctiveness values ($M = 0.5$, $SD = 0.0$), once the unique features were removed, and given that confounding psycholinguistic variables were accounted for in the ANCOVAs, the processing advantage for the distinctive features of nonliving compared with living things appears to be driven by the greater correlational strength of nonliving compared with living things. Clearly, in naturally occurring concepts, living and nonliving things will not be matched in this way, and we claim that both distinctiveness and correlation will contribute to their conceptual processing.

Comments on Cree et al. (2006): Feature and Concept Verifications of Distinctive Features

Cree, McNorgan, and McRae (2006) reported a series of experiments to determine whether distinctive features are processed more quickly than shared features and whether this difference interacts with domain (living vs. nonliving), as predicted by the CSA. Cree et al. paired a number of concepts with a distinctive feature (occurring on average in 1 concept in the feature norms) and a shared feature (occurring on average in 19 concepts in the feature norms) selected from their feature production norms (McRae et al., 2005). A portion of these concepts were entered into a computational model designed to simulate behavioral concept–feature verification performance, which generated the prediction that distinctive features would be processed more efficiently than shared features, irrespective of the domain to which the associated concepts belonged. This prediction contrasts with the domain by distinctiveness interaction predicted by the CSA, with worse performance being expected for the distinctive features of living things (Randall et al., 2004; Tyler & Moss, 2001). Cree et al. also carried out both a concept–feature verification experiment and a feature–concept verification study with their stimuli, which showed main effects of distinctiveness and stimulus onset asynchrony (SOA) but no interactions of distinctiveness with the living–nonliving domain. We note, however, that Cree et al.’s reported lack of domain effects must be viewed with caution, as comparisons of psycholinguistic variables of the shared and distinctive features were made within and not across domain, and potential differences in lexical and semantic factors in the stimuli (e.g., word and letter lengths of features, feature frequency, etc.) were not controlled for in their statistical analyses. These results were interpreted as supporting the authors’ feature-based conceptual attractor network model. Because the CSA predicts poorer performance with the distinctive features of living things compared with all other feature types (i.e., an interaction of distinctiveness with domain), these results were further interpreted as refuting the CSA.

In our view, Cree et al.’s (2006) findings do not contradict the CSA because the stimuli used in their experiments were not appropriate to test the model’s claim of an interaction of distinctiveness and correlation with domain. Cree et al. did not calculate the correlation values of their distinctive living and nonliving features, as they view these values as spurious (see above). Thus, these stimuli could not be used to test the central prediction of the CSA that the most distinctive features of nonliving things have a representational and processing advantage over the most distinc-

tive features of living things because of their higher correlation, because we do not know what the correlational statuses of these features were.

Nevertheless, Cree et al. (2006) raised important issues concerning the effects of distinctive features on conceptual processing. Their interpretation of the behavior of their computational model is particularly interesting. This model tested 13 living and 12 nonliving items from their concept production norms (although they do not specify the criteria used to select this subset from the set of 36 items used in the production norm study). Based on the ways in which the weights of the most distinctive and shared features developed over time, Cree et al. suggested that the semantic network first activates the most distinctive features, before forwarding activity to other related features and inhibiting activation from spreading to features that are shared with the target concept.

Finally, we note that Cree et al.’s (2006) distinctiveness effects were derived from a task suited to tap relatively late (and higher order) stages of conceptual processing: SOAs ranged from 300 to 1,650 ms, concept words were presented throughout the entire trial, and feature phrases were used instead of single words. For the Randall et al. (2004) behavioral study, on the other hand, we selected a paradigm specifically designed to tap into very early stages of conceptual processing by using a speeded feature verification task, that is, the visibility of the concept for 60 ms only, the use of a mask after concept presentation, the short (100-ms) presentation of the feature, and the fact that participants had to respond before a deadline beep. Thus, Randall et al.’s correlation effects in these early processing stages (reflected in the slower verification latencies to the distinctive living features) are compatible with McRae et al.’s (1997) previous simulations demonstrating faster rise times in the activation of more highly, as compared with less, correlated features. Taken together, these findings suggest that correlation may affect earlier stages and distinctiveness may influence later stages of conceptual processing.

We await future research that will further explore whether more or less distinctive features are facilitated as compared with shared features and determine whether and how these effects interact with feature correlation. We are currently investigating these issues with living and nonliving stimuli, which have a range of distinctiveness and correlation values and, critically, which are matched on conceptual structure over domain, using tasks designed to tap earlier and later stages of conceptual processing. We consider these experiments an even stronger test of the basic assumption of distributed models, namely, that conceptual structure, and not domain membership per se, determines conceptual organization and processing.

References

- Caramazza, A., & Mahon, B. Z. (2003). The organization of conceptual knowledge: The evidence from category-specific semantic deficits. *Trends on Cognitive Sciences*, 7, 325–374.
- Caramazza, A., & Mahon, B. Z. (2005). The organisation of conceptual knowledge in the brain: The future’s past and some future directions. *Cognitive Neuropsychology*, 22, 1–25.
- Caramazza, A., & Shelton, J. R. (1998). Domain-specific knowledge systems in the brain: The animate–inanimate distinction. *Journal of Cognitive Neuroscience*, 10, 1–34.
- Cree, G. S., McNorgan, C., & McRae, K. (2006). Distinctive features hold a privileged status in the computation of word meaning: Implications for

- theories of semantic memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 32, 643–658.
- Cree, G. S., & 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.
- Devlin, J. T., Gonnerman, L. M., Andersen, E. S., & Seidenberg, M. S. (1998). Category-specific semantic deficits in focal and widespread brain damage: A computational account. *Journal of Cognitive Neuroscience*, 10, 77–94.
- Durrant-Peatfield, M., Tyler, L. K., Moss, H. E., & Levy, J. (1997). The distinctiveness of form and function in category structure: A connectionist model. In M. G. Shafto & P. Langley (Eds.), *Proceedings of the nineteenth annual conference of the Cognitive Science Society* (pp. 193–198). Mahwah, NJ: Erlbaum.
- Gonnerman, L. M., Andersen, E. S., Devlin, J. T., Kempler, D., & Seidenberg, M. S. (1997). Double dissociation of semantic categories in Alzheimer's disease. *Brain and Language*, 57, 254–279.
- Greer, M., van Casteren, M., McClellan, S., Moss, H. E., Rodd, J., Rogers, T., & Tyler, L. K. (2001). The emergence of semantic categories from distributed featural representations. In J. D. Moore & K. Stenning (Eds.), *Proceedings of the 23rd annual conference of the Cognitive Science Society* (pp. 358–363). London: Erlbaum.
- Humphreys, G. W., & Forde, E. M. E. (2001). Hierarchies, similarity, and interactivity in object recognition: "Category-specific" neuropsychological deficits. *Behavioral and Brain Sciences*, 24, 453–509.
- Keil, F. (1987). Conceptual development and category structure. In U. Neisser (Ed.), *Concepts and conceptual development* (pp. 175–200). Cambridge, United Kingdom: Cambridge University Press.
- Malt, B. C., & Smith, E. (1984). Correlated properties in natural categories. *Journal of Verbal Learning and Verbal Behaviour*, 23, 250–269.
- Martin, A., & Chao, L. L. (2001). Semantic memory and the brain: Structure and processes. *Current Opinion in Neurobiology*, 11, 194–201.
- Martin, A., Ungerleider, L. G., & Haxby, J. V. (2000). Category-specificity and the brain: The sensory-motor model of semantic representations of objects. In M. S. Gazzaniga (Ed.), *The Cognitive Neurosciences* (2nd ed., pp. 1023–1036). Cambridge: MIT Press.
- Masson, M. (1995). A distributed memory model of semantic priming. *Journal of Experimental Psychology Learning, Memory, and Cognition*, 21, 3–23.
- McRae, K., & Cree, G. S. (2002). Factors underlying category-specific semantic deficits. In E. M. E. Forde & G. Humphreys (Eds.), *Category-specificity in mind and brain* (pp. 211–250). East Sussex, United Kingdom: Psychology Press.
- McRae, K., Cree, G. S., Seidenberg, M. S., & McNorgan, C. (2005). Semantic feature production norms for a large set of living and nonliving things. *Behavior Research Methods, Instruments, and Computers*, 37, 547–559.
- McRae, K., Cree, G. S., Westmacott, R., & de Sa, V. R. (1999). Further evidence for feature correlations in semantic memory. *Canadian Journal of Experimental Psychology*, 53, 360–373.
- McRae, K., de Sa, V. R., & Seidenberg, M. S. (1997). On the nature and scope of featural representations of word meaning. *Journal of Experimental Psychology: General*, 126, 99–130.
- Moss, H. E., Tyler, L. K., & Devlin, J. (2002). The emergence of category specific deficits in a distributed semantic system. In E. Forde & G. W. Humphreys (Eds.), *Category-specificity in brain and mind* (pp. 115–148). Sussex, United Kingdom: Psychology Press.
- Moss, H. E., Tyler, L. K., Durrant-Peatfield, M., & Bunn, E. M. (1998). 'Two eyes of a see-through': Impaired and intact semantic knowledge in a case of selective deficit for living things. *Neurocase*, 4, 291–310.
- Moss, H. E., Tyler, L. K., & Jennings, F. (1997). When leopards lose their spots: Knowledge of visual properties in category-specific deficits for living things. *Cognitive Neuropsychology*, 14, 901–950.
- Pexman, P. M., Holyk, G. G., & Monfils, M. H. (2003). Number-of-features effects and semantic processing. *Memory and Cognition*, 31, 842–855.
- Pexman, P. M., Lupker, S. J., & Hino, Y. (2002). The impact of feedback semantics in visual word recognition: Number-of-features effects in lexical decision and naming tasks. *Psychonomic Bulletin & Review*, 9, 542–549.
- Randall, B., Moss, H. E., Rodd, J. M., Greer, M., & Tyler, L. K. (2004). Distinctiveness and correlation in conceptual structure: Behavioral and computational studies. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 30, 393–406.
- Rosch, E. (1978). Principles of categorization. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and categorization* (pp. 27–48). Hillsdale, NJ: Erlbaum.
- Taylor, K. I., Moss, H. E., & Tyler, L. K. (2007). The conceptual structure account: A cognitive model of semantic memory and its neural instantiation. In J. Hart & M. Kraut (Eds.), *The neural basis of semantic memory* (pp. 265–301), Cambridge, United Kingdom: Cambridge University Press.
- Tyler, L. K., & Moss, H. E. (2001). Towards a distributed account of conceptual knowledge. *Trends in Cognitive Sciences*, 5, 244–252.
- Tyler, L. K., Moss, H. E., Durrant-Peatfield, M. R., & Levy, J. P. (2000). Conceptual structure and the structure of concepts: A distributed account of category-specific deficits. *Brain and Language*, 75, 195–231.
- Warrington, E. K., & McCarthy, R. (1983). Category-specific access dysphasia. *Brain*, 106, 859–878.
- Warrington, E. K., & McCarthy, R. (1987). Categories of knowledge: Further fractionations and an attempted integration. *Brain*, 110, 1273–1296.
- Warrington, E. K., & Shallice, T. (1984). Category-specific semantic impairment. *Brain*, 107, 829–853.

Received October 18, 2006

Revision received September 5, 2007

Accepted September 11, 2007 ■