Cognition 122 (2012) 363–374

Contents lists available at SciVerse ScienceDirect

Cognition

journal homepage: www.elsevier.com/locate/COGNIT

Contrasting effects of feature-based statistics on the categorisation and basic-level identification of visual objects

Kirsten I. Taylor [a,b,*], Barry J. Devereux [a], Kadia Acres [a,1], Billi Randall [a], Lorraine K. Tyler [a]

[a] Dept. of Experimental Psychology, University of Cambridge, United Kingdom
[b] Memory Clinic – Neuropsychology Center, University Hospital Basel, Switzerland

ARTICLE INFO

Article history:
Received 20 October 2010
Revised 1 November 2011
Accepted 2 November 2011
Available online 3 December 2011

Keywords:
Conceptual representation
Semantics
Object processing
Picture naming
Categorisation

ABSTRACT

Conceptual representations are at the heart of our mental lives, involved in every aspect of cognitive functioning. Despite their centrality, a long-standing debate persists as to how the meanings of concepts are represented and processed. Many accounts agree that the meanings of concrete concepts are represented by their individual features, but disagree about the importance of different feature-based variables: some views stress the importance of the information carried by distinctive features in conceptual processing, others the features which are shared over many concepts, and still others the extent to which features co-occur. We suggest that previously disparate theoretical positions and experimental findings can be unified by an account which claims that task demands determine how concepts are processed in addition to the effects of feature distinctiveness and co-occurrence. We tested these predictions in a basic-level naming task which relies on distinctive feature information (Experiment 1) and a domain decision task which relies on shared feature information (Experiment 2). Both used large-scale regression designs with the same visual objects, and mixed-effects models incorporating participant, session, stimulus-related and feature statistic variables to model the performance. We found that concepts with relatively more distinctive and more highly correlated distinctive relative to shared features facilitated basic-level naming latencies, while concepts with relatively more shared and more highly correlated shared relative to distinctive features speeded domain decisions. These findings demonstrate that the feature statistics of distinctiveness (shared vs. distinctive) and correlational strength, as well as the task demands, determine how concept meaning is processed in the conceptual system.

1. Introduction

Understanding how the meanings of concrete concepts are represented and processed stands at the heart of research on conceptual knowledge and has been approached from a number of different theoretical perspectives. Many models of conceptual knowledge assume some form of componentiality, where a concept is represented by its constituent parts, or features (Collins & Loftus, 1975; Murphy, 2002; Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976; Smith & Medin, 1981). One class of feature-based model claims that the statistical characteristics of features capture how concepts are represented. However, within this class of models there are disagreements about the functional role of statistical characteristics in accessing a concept's meaning (Cree, McNorgan, & McRae, 2006; Gonnerman, Andersen, Devlin, Kempler, & Seidenberg, 1997; McRae, de Sa, & Seidenberg, 1997; Randall, Moss, Rodd, Greer, & Tyler, 2004; Vinson, Vigliocco, Cappa, & Siri,

* Corresponding author. Address: Centre for Speech, Language and the Brain, University of Cambridge, Downing Street, Cambridge CB2 3EB, United Kingdom. Tel.: +44 (0)1223 765 938; fax: +44 (0)1223 766 452.

E-mail address: ktaylor@uhbs.ch (K.I. Taylor).

[1] Present address: Division of Psychology and Language Sciences, University College London, Chandler House, 2 Wakefield Street, London WC1N 1PF, United Kingdom.

0010-0277/$ - see front matter © 2011 Elsevier B.V. All rights reserved.
doi:10.1016/j.cognition.2011.11.001

2003). Here, we focus on two statistical characteristics of features which have played a key role in current theorising – the extent to which features co-occur in different concepts, and the extent to which features are distinctive to a particular concept or are shared by many concepts. We suggest that previously contradictory findings in the literature may be accounted for by the influence of task-dependent factors in determining how we understand concrete concepts.

The relevance of feature co-occurrence for conceptual representations was first highlighted by Rosch and colleagues (Rosch et al., 1976). They noticed that certain feature combinations frequently co-occurred, e.g. birds tended to have beaks, feathers and lay eggs. The existence of these feature clusters in natural categories led Rosch et al. (1976) to suggest that they had a special status in the conceptual representation of natural categories. Keil (1986) subsequently reported that clusters of co-occurring features were larger and more densely intercorrelated for living things than for nonliving things see also (Malt & Smith, 1984). The proposal that the co-occurrence of features differs across the two domains of knowledge (i.e. living vs. nonliving things) has since been supported by data from property norm studies, which collate features produced by healthy participants to a set of individual concepts. By calculating Pearson product-moment correlations across pairs of feature vectors derived from their property norm study of 190 concepts, McRae et al. (1997) observed that 11% of the feature pairs of living things were significantly correlated, compared to only 6% in nonliving thing concepts (see also Randall et al., 2004; Vinson et al., 2003).

While these studies showed that feature co-occurrence characterises the organisation of concepts, their functional relevance remained to be established. Early studies using category learning and typicality rating tasks found no effect of feature co-occurrence on performance (Malt & Smith, 1984; Murphy & Wisniewski, 1989). Instead, the effects of feature correlation were found when participants were required to explicitly compare correlated and uncorrelated feature pairs, and were presented with highly salient correlations (Malt & Smith, 1984). These studies appeared to indicate that feature co-occurrence was only functionally relevant when it was made explicit, or when participants were informed of the context for how features might co-occur (Murphy & Medin, 1985). However, McRae and colleagues (1997) explained the lack of feature co-occurrence effects as being due to the nature of the slow, off-line tasks used in these early studies, which required high-level reasoning processes (e.g. scripts, world knowledge). To test this hypothesis, McRae and colleagues contrasted performance on off-line concept similarity rating and typicality judgement tasks with that on speeded semantic priming and feature verification tasks. Consistent with their hypothesis, feature correlations predicted performance on the short SOA semantic priming and feature verification tasks, but not the concept similarity and typicality rating tasks (McRae et al., 1997). The facilitatory effect of feature correlation was subsequently replicated in a feature-concept priming study using lexical decision and single word stimuli to minimise syntactic processing: strongly intercorrelated features primed lexical decisions to target concepts significantly more than weakly intercorrelated features (Taylor, Moss, Randall, & Tyler, 2004). These and similar findings (Cree, McRae, & McNorgan, 1999; McRae, Cree, Westmacott, & de Sa, 1999; Randall et al., 2004) led to the proposal that strongly correlated features speed their activation in on-line comprehension tasks, a claim which has become a theoretical cornerstone in many current, feature-based accounts of conceptual representation and processing (Gonnerman et al., 1997; McRae, 2005; Moss, Tyler, & Taylor, 2007; Tyler & Moss, 2001; Vigliocco, Vinson, Lewis, & Garrett, 2004).

A second central variable in statistical feature-based accounts of conceptual knowledge is feature distinctiveness, or the extent to which features are distinctive to a particular concept or are shared by many concepts. The importance of distinctive features in the representation of concepts was already highlighted in the "classical view" of conceptual representations (Murphy, 2002; Smith & Medin, 1981). On this account, a concept could be classified as e.g. a tiger if it possessed all of the defining features of a tiger (Hull, 1920). Thus, defining features were sought that were cues to the identity of a concept. One key problem with this account is that it was not obvious how to identify defining features (Wittgenstein, 1953). To address this and other drawbacks (Medin & Smith, 1984; Murphy, 2002), a "probabilistic view" was developed that distanced itself from the defining feature view, and embraced the notion of feature similarity as a key representational dimension for concepts. According to this account, a feature belonged to a conceptual representation if it occurred in instances of the concept with a high probability, i.e. if it was shared across other instances of the concept (or category). Thus, an instance of a concept could be identified if it was sufficiently similar to a summary representation (Medin & Smith, 1984). As noted by Malt and Smith (1984), this model was recognised as incomplete since it neglected the importance of additional semantic information used during conceptual processing beyond shared features, e.g. correlated feature information.

Current statistical, feature-based views of conceptual representations unite the insights concerning defining and shared (similar) features into a single metric of distinctiveness: the inverse of the number of concepts a feature occurs in (Cree & McRae, 2003). Features that occur in many concepts are considered "shared" and have low distinctiveness values, whereas features that occur in very few concepts are considered "distinctive" and have high distinctiveness values. Thus, akin to defining features in the classical view, distinctive features are those that are true of a small number of concepts, and thus distinguish similar concepts from one another. For example, the distinctive feature ⟨has a hump⟩, but not the shared features ⟨has legs⟩ or ⟨has a tail⟩, distinguishes a camel from similar four-legged animals (Dean, Bub, & Masson, 2001).

Specific types of brain damage are associated with impairments in processing distinctive compared to shared features (Alathari, Trinh Ngo, & Dopkins, 2004; Martin, 1992; Moss & Tyler, 1997; Tyler & Moss, 2001; Warrington, 1975). Building on these patient findings, some investigators have hypothesised that distinctive features have a special representational status which confers a facilitatory effect during conceptual processing in healthy individuals (Cree et al., 2006; Hamilton & Geraci, 2006; Mirman & Magnuson, 2009). Cree et al. (2006) tested this hypothesis in concept-feature and feature-concept verification tasks in which features were either highly distinctive or shared. They found faster response times to highly distinctive compared to shared features, and suggested that following the presentation of either a concept or a feature, highly distinctive features are activated faster than shared features. Mirman and Magnuson (2009) extended these findings with a concept–concept similarity rating task, in which concepts sharing a highly distinctive feature were judged as significantly more semantically similar than concepts sharing a highly frequent (i.e. shared) feature. The results of these studies are consistent with a privileged status for distinctive features in facilitating conceptual processing.

Claims for the privileged status of distinctiveness are bolstered by neuropsychological findings demonstrating specific impairments for differentiating between highly similar concepts (Tyler & Moss, 2001). However, this behavioural impairment represents only part of the conceptual syndrome: patients' ability to understand other aspects of a concept's meaning, such as the category to which the concept belongs, is relatively spared (Moss, Tyler, & Devlin, 2002; Moss, Tyler, Durrant-Peatfield, & Bunn, 1998). These findings suggest that distinctive and shared properties play complementary, but integrated, roles in conceptual processing. Shared features are true of a high proportion of concepts in a category or domain, and thus are more informative about the category or domain to which a concept belongs. In contrast, distinctive features are true of a low proportion of concepts, and within-category differentiation therefore requires distinctive properties to distinguish between similar concepts. On this view, distinctive features facilitate the identification of instances of a concept at the basic level, while shared features facilitate category-level decisions, suggesting that depending on the task, different kinds of features will be differentially important (Moss et al., 1998, 2002; Taylor, Moss, & Tyler, 2007). This account is consistent with models of context-dependent automatic access of features during on-line conceptual processing (Potter & Faulconer, 1979; Tabossi, 1988), where the context in which a concept is encountered determines how it is processed or how it informs decision-making or response processes (see e.g. Grondin, Lupker, & McRae, 2009). Thus, a task-dependent account of the role of feature distinctiveness extends current models which assume a fixed pattern of featural activation or importance for any kind of conceptual decision.

As noted by Medin and Smith (1984), a limitation of probabilistic views of conceptual representation is that feature similarity (sharedness) is not the only variable relevant to conceptual processing; correlated feature information is also important. A key development in statistical, feature-based models of conceptual representations is the claim that correlational strength and distinctiveness interact to determine how concepts are processed, and the development of corresponding measures that combine correlational strength and distinctiveness information for individual features within a concept (Randall et al., 2004; Taylor et al., 2007; Tyler & Moss, 2001). One such measure is the extent to which different concepts' distinctive features are correlated. For example, living things are characterised by relatively few distinctive features which are weakly correlated, compared with nonliving things whose many distinctive features are relatively more strongly correlated (Moss et al., 2007; Randall et al., 2004; Taylor, Salamoura, Randall, Moss, & Tyler, 2008). Since correlational strength is predicted to facilitate conceptual processing (McRae et al., 1997, 1999; Taylor et al., 2004), and since the distinctive features of living things are more weakly correlated than the distinctive features of nonliving things (and the shared features of both living and nonliving things), Randall et al. (2004) predicted that healthy participants should be selectively disadvantaged at processing the distinctive features of living things. This prediction was confirmed in a speeded feature verification experiment designed to tap early and more automatic conceptual activation (Randall et al., 2004; Taylor et al., 2007). Thus, distinctiveness and correlation of individual features determine how conceptual information is processed in on-line tasks. Moreover, these claims can be integrated with that of a task-dependent utilisation of feature information to generate the prediction that basic-level identification will be modulated by the extent to which distinctive features are correlated, and category-level identification will be modulated by the extent to which shared features are correlated.

The aim of the present study was to test the hypothesis that the effects of distinctiveness and correlational strength of individual features in a concept determine the speed of conceptual processing as a function of the information required under different task conditions (task-dependent account; Potter & Faulconer, 1979; Tabossi, 1988; Taylor et al., 2007). The alternative hypothesis claims that correlated features, and distinctive features, facilitate conceptual performance, irrespective of task demands (the privileged status account; Cree et al., 2006; McRae et al., 1997). We developed a large set of concept stimuli for which we calculated extensive statistical feature information. The same set of stimuli was presented during two tasks with different informational requirements: a task in which participants were required to determine the basic-level identity of a visual concept, and a task in which participants were required to determine the domain (living vs. nonliving) to which the concept belonged. Both tasks were developed as regression designs to circumvent the drawbacks associated with factorial designs, i.e. small and potentially unrepresentative stimulus sets due to matching constraints. Data were analysed with mixed-effects models which simultaneously accounted for session and item (i.e. visual, lexical, phonologic and conceptual) and participant effects. Thus, we aimed to test the prediction that both correlational strength and distinctiveness determine conceptual processing in a task-specific manner by using a large, naturalistic set of stimuli and stringently testing for the effects of feature statistics on meaningful object processing while also accounting for the potential effects of a large number of participant, session, visual and psycholinguistic variables.

2. Materials and methods

Both experiments were developed as regression designs to circumvent the difficulty of matching stimuli in each condition of a factorial design on all the relevant psycholinguistic and perceptual variables, which may lead to relatively few unrepresentative stimuli in each condition. Moreover, both experiments used stimuli representing single concepts, since the presentation of single concepts has greater ecological validity compared with e.g. paradigms which present concept-feature pairs for verification decisions (Cree et al., 2006; Randall et al., 2004). Stimulus development included the following steps: (1) A large set of property norm data (McRae, Cree, Seidenberg, & McNorgan, 2005) formed the basis for stimulus selection. These property norm data contained a broad range of features belonging to each concept, including perceptual, functional, and encyclopaedic features. They had been collected from native North American English speakers, and were therefore first anglicised for use with native British English speaking participants. (2) The anglicised property norm data were used to calculate key feature statistic variables. (3) Pictures were obtained for appropriate concept stimuli. (4) We collated and collected information on other variables, such as visual and psycholinguistic variables (Dell, 1988; Levelt, 1999; Levelt, Roelofs, & Meyer, 1999). (5) A pre-test was run to acquire data on psycholinguistic and visual variables on each of the pictures in the set. (6) Finally, a PCA with varimax rotation was performed to reduce the number of "nuisance" (i.e., non-feature statistic) stimulus variables and the key feature statistic variables were residualised against the nuisance PCA component scores to ensure that each was orthogonal to the PCA components for use in the statistical analyses of the behavioural data. The factor loadings, residualised feature statistic variables, and relevant participant and session variables were used as predictors of reaction times (RTs) in mixed-effects models which simultaneously modelled item and subject random effects. These steps are described in detail below.

2.1. Anglicisation of property norms

Concepts were selected from a large property norm set (McRae et al., 2005). Since these norms were collected from North American English speakers, the concept and feature data were modified for use with native British English speakers in the following ways: We deleted 22 concepts and their corresponding features which are extremely unfamiliar to native British English speakers (e.g. "gopher"); 30 concept names were changed to their British English equivalents (e.g. "buggy" was changed to "pram"); features which were not appropriate in a British context were removed (e.g. mug – ⟨has a message⟩); in the original norms, features such as ⟨eats bread⟩ which describe two semantic properties (i.e. ⟨eats⟩ and ⟨eats bread⟩) were decomposed into their constituent features – for consistency, we decomposed features in the original norms which had not yet been separated, creating several new features. The resulting anglicised property norms included 517 concepts and 2159 features. Based on these norms, we calculated the number of features (NOF) for each concept in the norms. We consider NOF, which depends mainly on shared features, to be a measure of the semantic "richness" of concepts. Semantically rich concepts are suggested to benefit from a greater amount of activation at the conceptual level which cascades to facilitate, for example, phonological word form processing. Indeed, Pexman and colleagues showed that reading latencies, lexical decisions and semantic (concrete/abstract) decisions were facilitated for high compared to low NOF words (Pexman, Hargreaves, Siakaluk, Bodner, & Pope, 2008; Pexman, Holyk, & Monfils, 2003; Pexman, Lupker, & Hino, 2002).

2.2. Key feature statistic variables

Statistical analyses following McRae et al. (2005) and Randall et al. (2004) were conducted to produce the key theoretically interesting feature statistic variables. With respect to the statistics calculated over features, *feature distinctiveness* was calculated as 1/[number of concepts in which a feature occurred]. To measure the degree to which two features co-occur, we calculated each feature's *correlational strength*. This measure is based on Pearson's product moment correlations between pairs of features in the concept-feature matrix, where concepts are listed in rows and features are listed in columns, and each row $\times$ column value corresponds to the production frequency of a particular feature for a particular concept. The production frequencies reflected the number of participants providing the given feature in response to the corresponding concept (min. 5/30), excluding taxonomic features which reflect category-level, but not concept-specific, information. A feature's correlational strength was defined as the average of all significant correlation coefficients between the target feature and all other features in the norms. Since correlations with distinguishing features may be spurious (Cree et al., 2006; Taylor et al., 2008), mean correlational strength was calculated for shared features only (i.e., those occurring in more than two concepts).

Three key *concept*-specific variables were calculated from the feature-specific indices described above for use in the present analyses. The first was the 'mean distinctiveness' of all features belonging to a concept. The second was the average of all significant pairwise correlations between the shared features of a concept ('correlational strength'). A final feature-based statistic variable was developed to estimate the relationship between the correlational strength of a concept's shared and more distinctive features. In a distributed conceptual system, these relative (as opposed to absolute) differences in feature statistics are assumed to be responsible for behavioural effects in normal and brain-damaged systems (Taylor et al., 2008; Tyler, Moss, Durrant-Peatfield, & Levy, 2000). To estimate this relationship, we created a 'correlation $\times$ distinctiveness' measure, which was the unstandardised slope of the regression line describing the scattergraph of each concept's features, where correlational strength was on the *x*-axis and feature distinctiveness on the *y*-axis. This measure therefore reflects the relative correlational strengths of the continuum of feature distinctiveness values, consistent with graded properties of a distributed, feature-based system. Only shared features (i.e. distinctiveness values <.5) were used to generate the 'correlation × distinctiveness' values as it has been argued that the correlational strengths of distinguishing features are spurious (Cree et al., 2006; Taylor et al., 2008). Thus, living things typified by many shared, highly correlated features, and relatively few, weakly correlated distinctive features have a low 'correlation × distinctiveness' value ($M = 0.61$, $SD = 0.79$, see Fig. 1a). Nonliving things whose distinctive features are relatively more highly correlated have higher 'correlation × distinctiveness' values ($M = 1.05$, $SD = 0.94$; $t(515) = 5.445$, $p < .0001$; see Fig. 1b).

2.3. Picture selection

From the set of 517 anglicised concepts words, we excluded homonyms and concepts which could not be visualised independent of their context (e.g. "flat"). Naturalistic pictures were found for a total of 412 concepts. We deleted contextual information from all pictures so that each object appeared against a white background and resized images to fit comfortably on a computer monitor with the longest picture axis set to a width of 750 pixels or height of 550 pixels.

2.4. Non-feature statistic stimulus variables

Picture categorisation and basic-level naming are known to be influenced by visual, phonological, lexical and semantic factors in addition to the potential effects of conceptual structure (Glaser, 1992; Johnson, Paivio, & Clark, 1996; Stewart, Parkin, & Hunkin, 1992). To control for the effects of these variables while testing for the conceptual structure effects of interest, we collated and collected a large number of concept-specific variables for inclusion in our mixed model behavioural analyses. These variables and their effects on picture categorisation and basic-level naming are described in Supplementary methods. All word form variables were calculated based on the most common correct name produced from the overt basic-level naming Experiment 1.

2.5. Stimuli pre-test

A pre-test was conducted to obtain data on psycholinguistic and visual variables on each picture. This pre-test was conducted with an independent group of 17 healthy participants who did not participate in Experiments 1 and 2. Participants were first presented with the 412 pictures (and 12 filler pictures) for familiarity ratings (7-point scale, where 1 = completely unfamiliar). Following the familiarity ratings, the same participants viewed the pictures a second time, in a different order. This time, pictures were presented together with their verbal label. Participants were instructed to rate the exemplarity of each picture as well as the picture's visual complexity (7-point scales, where 1 = *poor picture of concept word* and *visually very simple*, respectively). For each picture, familiarity, exemplarity and visual complexity ratings were averaged over all participants.

2.6. Principal components analysis

Many of the visual, phonological, lexical and semantic variables described above are highly intercorrelated. To reduce the number of covariates in our analyses of the effects of feature statistic variables on meaningful object processing, while taking into account their potential interdependencies, we performed a principal components analysis (PCA) with varimax rotation on the non-feature statistic variables describing the object pictures.

Where necessary, variables were transformed to achieve a more optimal normal distribution (transformations noted in Table 1). PCA analyses with varimax rotation produced an eight-factor solution which accounted for 85.5% of the variance and provided the most meaningful set of components (see Table 1). The components were interpreted as: (1) Phonological, (2) Visual Complexity, (3) Naming Agreement, (4) Frequency, (5) Familiarity, (6) Cohort Size, (7) H-statistic, and (8) NOF. Of note is the lack of a heavy loading for the rated exemplarity measure. The three key feature statistic variables (mean distinctiveness, correlational strength of shared features, and 'correlation × distinctiveness') were then residualised against the eight components to ensure their orthogonality to the components. The eight

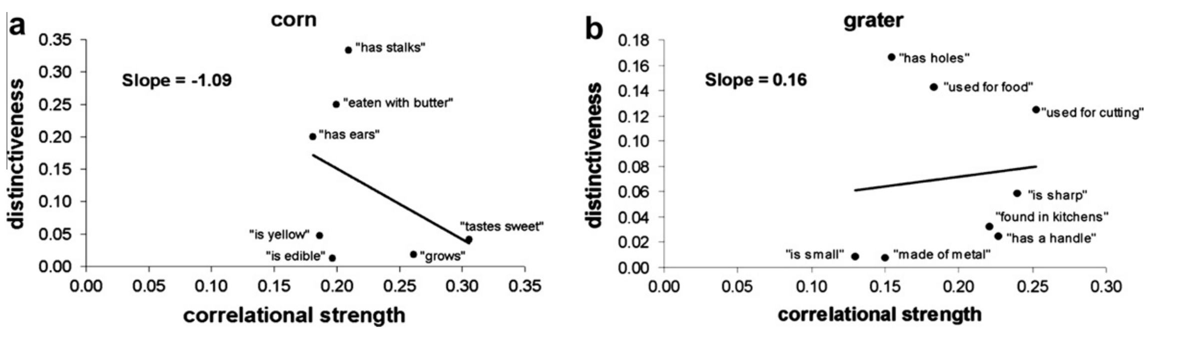

Fig. 1. Examples of the correlation × distinctiveness variable reflecting the relationship between correlational strength status of features with different distinctiveness values: (a) the more distinctive features in the concept "corn" have a weaker correlational strength than the shared features of corn, which are more strongly correlated, resulting in a negative value for 'correlation × distinctiveness'; (b) both the more and less distinctive features of "grater" are similarly highly correlated, resulting in a value of 'correlation × distinctiveness' around zero.

Table 1
Results of the principal components analysis with varimax rotation (transformation noted parenthetically) of the nuisance variables (factor loadings exceeding .75 are bolded).

	Component							
	1 Phonology	2 Visual complexity	3 Name Agreement	4 Frequency	5 Familiarity	6 Cohort size	7 H-stat	8 NOF
Phonological neighbourhood density (sqrt)	**.946**	–.030	.014	.190	–.006	–.044	.003	.050
Onset phonological neighbourhood density (sqrt)	**.921**	–.033	–.008	.110	.020	.131	.027	.046
Number of phonemes	**–.837**	.040	.001	–.267	.017	.208	–.061	.001
Lemma frequency (ln(x + 1))	.287	.044	.153	**.897**	.120	–.050	–.012	.092
Word form frequency (ln(x + 1))	.313	.035	.160	**.891**	.111	–.007	.015	.089
Number of features (NOF)	.083	–.008	.082	.135	.118	.063	–.027	**.948**
Visual complexity – number of pixels in picture	–.087	**.777**	–.016	.225	.115	–.104	.218	.094
Visual complexity – filesize	–.062	**.936**	–.014	.051	–.032	–.055	.092	.037
Visual complexity – bits per pixel	.071	**.663**	–.022	–.046	.229	.180	–.131	–.175
Visual complexity – subjective ratings	–.031	**.665**	–.002	–.190	–.384	–.087	–.094	.022
% Naming agreement	–.005	–.043	**.940**	.145	.118	.002	–.189	.058
% Concept agreement	.016	–.022	**.955**	.139	.121	–.011	.121	.034
Familiarity	.026	–.054	.071	.292	**.868**	.077	.052	.042
Exemplarity	–.055	.181	.351	–.127	**.659**	–.226	–.111	.169
H-statistic	.070	.072	–.050	–.002	–.007	.039	**.963**	–.026
Cohort size	–.070	–.038	–.012	–.042	–.032	**.952**	.037	.061

PCA components and three residualised feature statistic variables were used as independent predictor variables in the statistical analyses. The results regarding the PCA components, which are not the focus of theoretical interest, are described in Supplementary material.

3. Experiment 1

In Experiment 1 participants were presented with the 412 pictures described above for basic-level naming, a task which requires distinctive features in order to differentiate between similar concepts. We therefore predict that performance on this task will be facilitated in concepts with relatively more distinctive features. We further predict that the relative correlational strengths of shared and distinctive features will influence performance, such that basic-level identification is facilitated in concepts with more highly correlated distinctive relative to shared features (i.e., a facilitatory effect of 'correlation × distinctiveness'). We note that 'correlation × distinctiveness' differentiates living things from nonliving things, since the former are characterised by weakly correlated distinctive features and the latter by strongly correlated distinctive features, although this domain was not manipulated here. Critically, these predictions for basic-level identification contrast with those for domain-level identification (Experiment 2): since the latter relies on shared feature information, we predict that responses in this condition will be facilitated by relatively more shared and relatively more highly correlated shared compared to distinctive features.

3.1. Methods

3.1.1. Participants

A separate group of 20 healthy native British-speaking individuals from the Centre for Speech, Language and the Brain (CSLB) volunteer panel participated in Experiment 1 (13 women; mean age 26.1 y (range: 18–41 y)). All participants had normal or corrected-to-normal vision, none were colour blind and they were paid for their participation.

3.1.2. Procedure

The 412 picture stimuli were presented to participants in the same, pseudorandomised order in two blocks each comprising two lead-in and 206 target items. Each trial consisted of a picture presented for 2000 msec followed by a blank screen for 600 msec. Participants were seated in front of the computer monitor and fitted with a headset with microphone which was placed close to their mouth. The sensitivity of the voice key trigger was adjusted to each participant. Participants were instructed to overtly name the pictured object at the basic-level as quickly and accurately as possible, and to suppress extraneous speech noises such as coughing, throat clearing and lip smacking. Following a practice session with an independent set of eight stimuli, the voice key trigger was readjusted to account for the reduction in participants' speech response volume during the session. DMDX software controlled stimulus presentation and response collection (Forster & Forster, 2003).

3.1.3. Statistical analyses

Since the purpose of the experiment was to determine the influence of specific feature statistic variables on basic-level naming, and to ensure that the lexical, phonological and cohort measures reflected the word form retrieved by the participants, we were only able to include stimuli in the statistical analyses which were correctly conceptually identified by over 70% of participants. For the 302 stimuli which fulfilled this criterion, we excluded incorrect responses, responses which involved stuttering or were faster than 300 msec, and timeouts greater than 4 s (28% of the data).

The remaining RT data were inverse transformed (Ulrich & Miller, 1994) and analysed with mixed-effects models in R (R Development Core Team, 2008) with the lmer (Pinheiro & Bates, 2000) and languageR (Baayen, 2008) libraries. Mixed-effects models use the full dataset from all participants to model items and participants as crossed random effects and determine the fixed effects of predictor variables of interest (i.e. the eight PCA components and three residualised feature statistic variables). In addition, the variance associated with participant and session variables can also be entered into the model, thereby explaining more of the data's variance and producing a more valid estimate of the effects of the predictors (Baayen, Davidson, & Bates, 2008). The mixed-effects model therefore contained the presentation number and the inverse of the preceding trial's response time as session-specific predictor variables (these effects are not commented on below as they are of no theoretical interest) as well as the predictors of interest. We report *p*-values resulting from 200,000 samples from the posterior distribution of the parameters of the fitted model using Markov Chain Monte Carlo methods (Baayen, 2008).

3.2. Results and discussion

The harmonic mean naming latency across all items and participants was 868 ms (overall SD = 218 ms). The results of the mixed effects analyses with the eight PCA components and three residualised feature statistic variables are shown in Table 2. The results of the non-critical PCA components are described in Supplementary material.

Although current models claim that correlational strength *per se* speeds conceptual processing (McRae et al., 1997, 1999; Tyler & Moss, 2001), we found that the correlational strength of shared features did not affect naming RTs ($t = -0.182$, $p = .857$). Basic-level naming latencies were however significantly facilitated for concepts with more distinctive features (mean distinctiveness; $t = 1.923$, $p = .04$). This finding is consistent with claims for the primacy of distinctive feature effects in any type of conceptual processing (Cree et al., 2006), as well as the view that distinctive features are especially important for differentiating between highly similar concepts (Moss et al., 2002; Tyler & Moss, 2001). Finally, the 'correlation × distinctiveness' component significantly facilitated naming RTs ($t = 3.226$, $p < .001$), with faster responses to concepts with relatively more strongly correlated distinctive relative to shared features. This finding supports the view that the relative correlational strengths of a concept's shared and more distinctive properties influences conceptual processing, i.e. that the relatively greater correlational strengths of distinctive features critical for basic-level identification (as indexed here by naming latencies) speeds their on-line activation when objects must be identified at the basic-level (Randall et al., 2004; Taylor et al., 2007; Tyler & Moss, 2001).

4. Experiment 2

Experiment 2 was designed to test the predictions (see Introduction) that shared features, and relatively highly correlated shared features, facilitate domain-level decisions which rely on common (shared) feature information. These predictions are diametrically opposed to the predicted effects in basic-level naming tasks which rely on distinctive and relatively highly correlated distinctive feature information.

4.1. Methods

4.1.1. Participants

Twenty healthy, right-handed, native British-speaking individuals from the CSLB volunteer panel, who had not taken part in Experiment 1, participated in Experiment 2 (10 women; mean age 19.8 y (range: 18–31 y)). All participants had normal or corrected-to-normal vision, none were colour blind, and all were paid for their participation.

4.1.2. Procedure

The set of 412 picture stimuli included 182 items categorised as living things and 235 as nonliving things. To balance the numbers of living and nonliving concepts, we supplemented the target picture set with an additional 58 pictures of living thing stimuli which were not included in the data analyses. Stimuli were presented in the same, pseudorandomised order to all participants in two blocks of 235 stimuli each. Presentation timing was identical to Experiment 1 to maximise comparability over studies.

Table 2
Results of the mixed-effects models for Experiments 1 and 2 demonstrating effects of each PCA component and the residualised feature statistic variables (bolded) on inverse transformed reaction times.

Variable	Overt naming				Domain decision			
	Estimate	Standard error	*t*-value	*p*-value	Estimate	Standard error	*t*-value	*p*-value
Phonology	0.005	0.006	0.782	0.416	−0.007	0.009	−0.777	0.408
Visual complexity	0.000	0.006	−0.067	0.925	0.014	0.009	1.547	0.083
Naming agreement	0.123	0.020	6.020	0.000	−0.012	0.009	−1.342	0.141
Frequency	0.051	0.006	7.890	0.000	−0.020	0.009	−2.144	0.019
Familiarity	0.079	0.007	11.924	0.000	0.015	0.009	1.698	0.056
Cohort size	−0.015	0.007	−2.236	0.014	−0.020	0.009	−2.323	0.010
H-statistic	−0.020	0.008	−2.504	0.009	−0.011	0.009	−1.220	0.170
NOF	0.026	0.007	4.047	0.000	0.029	0.009	3.310	0.000
Correlational strength	**−0.001**	**0.007**	**−0.182**	**0.857**	**0.014**	**0.009**	**1.532**	**0.095**
Mean distinctiveness	**0.013**	**0.007**	**1.923**	**0.040**	**−0.027**	**0.010**	**−2.771**	**0.002**
'Correlation × distinctiveness'	**0.021**	**0.006**	**3.226**	**0.000**	**−0.020**	**0.009**	**−2.230**	**0.014**

Thus, in each trial, a picture appeared in the middle of a computer screen for 2000 msec, followed by a blank computer screen for 600 msec. Participants were instructed to decide whether each picture depicted a living or a nonliving thing, and to press a corresponding button on a button box as quickly and as accurately as possible. Half of the participants responded "living" with their right and "nonliving" with their left hand, and the other half of participants received the reverse hand-response assignment. The experiment was preceded by a short practice session with a new set of five picture stimuli. DMDX software controlled stimulus presentation and response collection (Forster & Forster, 2003).

4.1.3. Statistical analyses

RT data from correct responses to the 412 target items were inverse transformed prior to analyses (Ulrich & Miller, 1994) and were analysed with mixed-effects models in R (R Development Core Team, 2008) with the lmer (Pinheiro & Bates, 2000) and languageR (Baayen, 2008) libraries (see Experiment 1). Besides the eight PCA components and three residualised feature statistic variables, the following participant and session variables were included in the model but are not reported below because they are of no theoretical interest: The response hand assignment, presentation number, the participant's response (living vs. nonliving), whether the same button was pressed on the preceding trial, whether the previous response was an error, whether the previous trial presented an object picture from the same domain, and the inverse of the preceding trial's response time. We report p-values resulting from 200,000 samples from the posterior distribution of the parameters of the fitted model using Markov Chain Monte Carlo methods (Baayen, 2008).

4.2. Results and discussion

Participants performed the task accurately (97.2% correct), with a harmonic mean RT of 539 ms (overall SD = 179 ms). The results of the mixed effects analyses with the eight PCA components and three conceptual structure variables are shown in Table 2, and the results of the non-critical PCA components are described in Supplementary materials.

As predicted by the task-dependent, but not privileged status, account of feature distinctiveness, domain decision RTs were significantly faster to concepts with many *shared* relative to distinctive features (mean distinctiveness; $t = -2.771$, $p < .01$). Moreover, concepts with relatively highly correlated shared (relative to distinctive) features significantly facilitated domain decision RTs ($t = -2.230$, $p = .01$): the more strongly correlated a concept's shared features were (relative to its distinctive features), the faster the domain decision RTs. Thus, these findings indicate that domain decisions rely heavily on shared features which indicate category membership, such that the number of shared features and a greater relative correlational strength of these features significantly facilitate domain decisions. Taken together with the findings from Experiment 1 in which concepts with relatively more distinctive than shared features and relatively more highly correlated distinctive features facilitated basic-level identification latencies, these results support a task-dependent account of feature statistics on conceptual processing.

5. General discussion

Basic-level naming latencies were significantly facilitated for concepts with more distinctive relative to shared features, and relatively more strongly correlated distinctive features. In contrast, categorisation decisions were significantly facilitated for concepts with more shared relative to distinctive features, and with relatively more strongly correlated shared features. These findings support the validity of distributed, feature-based accounts of conceptual knowledge, and the cognitive relevance of feature-based statistics in object recognition. Importantly, relative, concept-specific variables best predicted performance, indicating that conceptual processing is driven by the relative statistical characteristics of a concept's features. The present, complementary set of findings was obtained by analysing the data collected on the same stimuli with mixed-effects models which simultaneously controlled for the influence of numerous participant, session, visual, phonologic, lexical and other semantic variables. Moreover, all experiments involved classifying or identifying single instances of concepts, thereby avoiding tasks which make explicit the relationship between concepts and features.

5.1. Effects of feature statistics

The present findings allow us to adjudicate between different accounts of the effects of feature statistic variables on semantic processing. It is generally agreed that correlational strength, i.e. the degree to which two features co-occur, speeds conceptual processing. Correlational strength between features is thought to be accrued through repeated exposure and thus co-activation of the cluster of features (e.g. ⟨has legs⟩, ⟨has a tail⟩ and ⟨has eyes⟩). Within a Hebbian framework, co-activation of cell assemblies strengthens their interconnection (Hebb, 1949), thereby purportedly making correlated features more resilient to the effects of brain damage (Gonnerman et al., 1997). Strongly correlated features are hypothesised to benefit from mutual co-activation to speed their automatic activation in on-line comprehension tasks (McRae et al., 1997). In contrast, there are different views on the role of distinctiveness on on-line conceptual processing. On some models, distinctive features have a special representational status which functions to facilitate conceptual processing in healthy individuals (Cree et al., 2006; Hamilton & Geraci, 2006; Mirman & Magnuson, 2009). Alternatively, task-dependent accounts predict that the task context determines which features are important for the conceptual task at hand: since distinctive features are true of a small proportion of concepts, they facilitate basic-level identification: since shared features are true of a large proportion of concepts in a specific category or domain, they facilitate category-level (or domain-level) decisions. We suggest that the combined effects of correlational strength and distinctiveness influence conceptual processing, such that relatively more correlated distinctive features speed basic-level identification, and relatively more highly correlated shared features facilitate category-level decisions.

As predicted by both the privileged status and task-dependent accounts, basic-level naming latencies were significantly facilitated by high mean distinctiveness values, i.e. for concepts with relatively more distinctive than shared features. Naming latencies were also faster to concepts with relatively more highly correlated distinctive features ('correlation $\times$ distinctiveness' variable). This measure reflects the relationship between the correlational strengths of a concept's shared and distinctive features. Thus, it has the advantage that it assesses the relative correlational strengths of the continuum of feature distinctiveness values, consistent with graded properties of a distributed, feature-based system, and excludes correlational strengths of distinguishing features which may be spurious (Cree et al., 2006; Taylor et al., 2008). A unique and central prediction of accounts which stress the multidimensional nature of feature statistics is that the correlational strength of relatively more distinctive and shared features exerts functional effects on conceptual processing (Moss et al., 2002; Randall et al., 2004; Taylor et al., 2007, 2008), a prediction confirmed by the present experiment.

The results of the domain-level identification experiment revealed a complementary pattern of results. Domain decision latencies were significantly facilitated for concepts with more highly correlated shared relative to distinctive features, in line with a facilitatory role of correlational strength (McRae et al., 1997). Domain decisions were also facilitated for concepts with relatively more shared than distinctive features. The relative importance of shared over distinctive features in domain decisions is consistent with the task-dependent account of feature distinctiveness, but not with claims for the privileged role of distinctive features in conceptual processing (Cree et al., 2006; Mirman & Magnuson, 2009). Instead, both sets of findings support a feature-based conceptual architecture in which different aspects of meaning representations are flexibly activated during conceptual processing and/or differentially inform decision-making processes as a function of task demands and conceptual feature statistics.

5.2. Feature statistics and object recognition

We claim that visual information interfaces with the semantic system comprised of features which are characterised by the key statistical properties of correlational strength and distinctiveness. The activation of a particular concept is driven by the relative proportion of shared and distinctive features and the relative correlational strengths of a concept's shared and more distinctive features (Randall et al., 2004; Taylor et al., 2007, 2008; Tyler & Moss, 2001; Tyler et al., 2000). Task may function to influence the late response stage by e.g. reading out the activation level of features critical for the task, or by influencing the initial rise time in feature activation. This latter hypothesis is supported by the finding that task demands affected the feature statistic variable based on correlational strength (correlational strength $\times$ distinctiveness) known to affect early rise time of feature activation (McRae et al., 1997). Once activity has settled onto a unique conceptual representation, this information is output to the verbal or motor system for the appropriate response. Thus, this account predicts task-dependent effects of distinctiveness/sharedness and correlational strength of shared and more distinctive features.

The Hierarchical Interactive Theory (HIT) of visual object recognition makes similar predictions regarding task-dependent effects in picture naming (Humphreys & Forde, 2001; Humphreys, Price, & Riddoch, 1999). This model proposes that object processing is interactive and comprises the processing of structural (visual) descriptions of objects, from which semantic representations are accessed, followed by the access of phonological codes in the case of basic-level naming (Humphreys & Forde, 2001; Humphreys et al., 1999). The similarity of structural descriptions is claimed to prominently influence both categorisation and naming processes, but in different ways. An object which shares structural descriptions with many other category members (e.g. where there is a high degree of feature overlap in visual features, as is the case for living things) will produce a large amount of overlapping activation in the structural descriptions of other category members, providing positive evidence of category membership which speeds categorisation latencies (Price & Humphreys, 1989; Snodgrass & McCullough, 1986) and domain decisions (Riddoch & Humphreys, 1987). In contrast, the similarity of structural descriptions slows basic-level naming latencies because additional, e.g. semantic information is required to differentiate the object from structurally similar competitors (Humphreys & Forde, 2001; Humphreys et al., 1999; Price & Humphreys, 1989). Thus, the present findings of complementary effects of distinctiveness/sharedness on basic-level naming and domain categorisation tasks, respectively, are consistent with the HIT's predictions (Humphreys & Forde, 2001; Humphreys et al., 1999).

The HIT model was developed to account for the effect of visual features (structural descriptions) and semantic features on *visual* object recognition. Therefore, other kinds of sensory information and non-perceptual (e.g. functional) features are not considered by the HIT. Moreover, the HIT does not account for the effects of correlational strength of features. In contrast, a task-dependent view of the effects of feature statistics does not differentiate between representational stores for visual and semantic information, but claims that all meaningful features are coded in the same feature-based conceptual space. Thus, the effects of not only visual, but also other sensory and non-perceptual (e.g. functional) features can be accounted for. Moreover, based on the same architecture, the task-dependent view generates additional predictions regarding the effects of feature co-occurrence (correlational strength). We postulate a single, feature-based conceptual system in which conceptual processing depends on the combined effects of the statistical characteristics of the relevant feature in the context of the task, such that no additional processes or information are required to explain conceptual performance on different kinds of tasks.

A parallel distributed processing account of semantic memory similar in architecture to the HIT has been proposed by Rogers and colleagues (2004). This model contains three layers coding for (1) visual object features, (2) hidden semantic units and (3) verbal information including name units (basic-level or category names), and verbal perceptual, functional and encyclopaedic feature description units. Data from visual and verbal property generation studies were used to construct visual and verbal feature vectors for 48 concepts. These feature vectors were encoded or manipulated to ensure that they produced concept-specific and category-congruent patterns of activity. For example, distinctive (idiosyncratic) features were manipulated to ensure that individual concepts produced unique patterns of activity, and feature vectors were coded to specify whether the feature was observed in the category, idiosyncratic (not always observed in the category), or never appeared in the category, and distortions were made to the artefact/fruit feature vectors to render them more visually similar to artefacts than animals and more verbally distinguishable from artefacts and animals. Thus, this model relies on feature distinctiveness (whether a feature is shared within or distinctive to (idiosyncratic) a category of concepts), as well as investigator adjustments to maximise conceptual differentiation within and between categories. When lesioned, the trained model accounted for a variety of behavioural data from brain-damaged patients, in particular those with semantic dementia (Rogers et al., 2004).

We claim that conceptual representations and category or domain structure emerges out of the distributed feature system coding solely for feature distinctiveness/sharedness and correlational strength, and thus do not impose additional constraints to differentiate concepts, categories or domains (Moss et al., 2002; Tyler & Moss, 2001). Thus, we calculated feature statistic data across a large set of 541 concepts without explicitly coding for category or domain information. The present results show that these feature statistic variables suffice to explain object identification and domain decisions in normal conceptual processing, as well as the behavioural performance of brain-damaged individuals (described below), without recourse to additional concept or categorical constraints.

5.3. Feature statistics and domain dissociations

Statistical, feature-based accounts of conceptual representations claim that feature statistics determine domain dissociations, most remarkably observed in patients with category-specific semantic impairments for living things (Gonnerman et al., 1997; McRae, 2005; Moss et al., 2002, 2007; Taylor et al., 2007; Tyler & Moss, 2001). Correlational strength is hypothesised to protect features from the effects of brain damage, and thus will especially protect the shared features of living and nonliving things, and the distinctive features of nonliving things, which are typically highly correlated (Moss, Rodd, Stamatakis, Bright, & Tyler, 2005; Randall et al., 2004). We therefore predict that shared features (irrespective of domain) will be relatively preserved even in patients with category-specific semantic impairments for living things, as will the distinctive features of nonliving things, while the distinctive features of living things will be most severely affected. Thus, patients will be especially impaired on tasks which require access to the distinctive properties of living things, but exhibit relatively normal performance on tasks requiring shared feature information. These predictions have been previously tested in studies which vary task demands. For example, Moss and colleagues demonstrated that while such patients were unable to uniquely identify living things, they were able to sort concepts according to superordinate category (Moss et al., 1998), produce shared, highly correlated concept features in drawing tasks (Moss, Tyler, & Jennings, 1997), and verify the highly correlated, shared properties of both living and nonliving things (Moss et al., 1998; see also Moss et al., 2002). Extending these patient findings, the present experiments demonstrate that feature statistics significantly influences unique identification and categorisation in healthy subjects in a mirror-image fashion.

6. Conclusions

The present findings demonstrate that the key feature statistic variables of feature distinctiveness and correlational strength determine how meaningful visual objects are processed within the context of task demands. The findings provide strong evidence for statistical, feature-based models of meaning (McRae et al., 1997; Tyler & Moss, 2001; Vigliocco et al., 2004). Moreover, these effects unite previously divergent views on the precedence of concept-specific, defining (distinctive) features (classical view; privileged access account) and common, shared features (probabilistic view) into a single account. While these effects were demonstrated with visual object concepts, we assume that they reflect modality independent processing principles of feature statistic, and would therefore be observable irrespective of input modality. In the present experiments, we attempted to move beyond a one-dimensional approach to feature statistics which explores the effects of each feature statistic variable in isolation. Towards this end, we developed relational measures of statistical feature properties, e.g. the relative correlational strengths of shared vs. more distinctive features within a concept (correlation $\times$ distinctiveness). These variables reflect the view that conceptual processing corresponds to the activation of bundles of features within a concept whose relative statistical properties determine conceptual processing, and that the activation of features is flexible and depends in part on the kinds of conceptual information required to perform the task at hand.

Acknowledgements

This study was supported by grants from the British Academy (LRG-45583), Newton Trust, by funding from the European Research Council under the European Community's Seventh Framework Programme (FP7/2007-2013)/ERC Grant agreement no. 249640 (LKT) and by a Marie Curie Intra-European Fellowship and Swiss National Science Foundation Ambizione Fellowship (KIT). We thank Dr. Ken McRae for kindly providing us with the production frequency vectors necessary to calculate new feature statistic statistics for the anglicised feature production norms. We also thank Marie Dixon for her help in scoring the basic-level naming responses.

Appendix A. Supplementary material

Supplementary data associated with this article can be found, in the online version, at doi:10.1016/j.cognition.2011.11.001.

References

Alathari, L., Trinh Ngo, C., & Dopkins, S. (2004). Loss of distinctive features and a broader pattern of priming in Alzheimer's disease. *Neuropsychology, 18*(4), 603–612. doi:10.1037/0894-4105.18.4.603.

Baayen, R. (2008). *Analyzing linguistic data. A practical introduction to statistics using R*. Cambridge, UK: Cambridge University Press.

Baayen, R., Davidson, D., & Bates, D. (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language, 59*, 390–412.

Collins, A., & Loftus, E. (1975). A spreading-activation theory of semantic processing. *Psychological Review, 82*, 407–428.

Cree, G., 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., & 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.

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

Dean, M., Bub, D., & Masson, M. (2001). Interference from related items in object identification. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 27*(3), 733–743.

Dell, G. (1988). The retrieval of phonological forms in production: Test of predictions from a connectionist model. *Journal of Memory & Language, 27*, 124–142.

Forster, K., & Forster, J. (2003). A windows display program with millisecond accuracy. *Behavior Research Methods, Instruments, and Computers, 35*, 116–124.

Glaser, W. (1992). Picture naming. *Cognition, 42*, 61–105.

Gonnerman, L., Andersen, E., Devlin, J., Kempler, D., & Seidenberg, M. (1997). Double dissociation of semantic categories in Alzheimer's disease. *Brain and Language, 57*, 254–279.

Grondin, R., Lupker, S. J., & McRae, K. (2009). Shared features dominate semantic richness effects for concrete concepts. *Journal of Memory and Language, 60*(1), 1–19.

Hamilton, M., & Geraci, L. (2006). The picture superiority effect in conceptual implicit memory: A conceptual distinctiveness hypothesis. *The American Journal of Psychology, 119*(1), 1–20.

Hebb, D. O. (1949). *The organization of behavior: A neuropsychological theory*. New York: Wiley.

Hull, C. L. (1920). Quantitative aspects of the evolution of concepts. *Psychological Monographs, 28*, 1–85.

Humphreys, G., & Forde, E. (2001). Hierarchies, similarity, and interactivity in object recognition: "Category-specific" neuropsychological deficits. *Behavioral and Brain Sciences, 24*, 453–509.

Humphreys, G., Price, C., & Riddoch, M. (1999). From objects to names: A cognitive neuroscience approach. *Psychological Research, 62*, 118–130.

Johnson, C., Paivio, A., & Clark, J. (1996). Cognitive components of picture naming. *Psychological Bulletin, 120*, 113–139.

Keil, F. (1986). The acquisition of natural kinds and artifact terms. In W. Demoupoulous & A. Marras (Eds.), *Language learning and concept acquisition: Foundational issues* (pp. 133–153). Norwood, NJ: Ablex.

Levelt, W. (1999). Models of word production. *Trends in Cognitive Sciences, 3*, 223–232.

Levelt, W., Roelofs, A., & Meyer, A. (1999). A theory of lexical access in speech production. *Behavioral and Brain Sciences, 22*, 1–75.

Malt, B. C., & Smith, E. E. (1984). Correlated properties in natural categories. *Journal of Verbal Learning and Verbal Behaviour, 23*, 250–269.

Martin, A. (1992). Degraded knowledge representations in patients with Alzheimer's disease: Implications for models of semantic and repetition priming. In L. Squire & N. Butters (Eds.), *Neuropsychology of memory* (pp. 220–232). New York: Guilford Press.

McRae, K. (2005). Semantic memory: Some insights from feature-based connectionist attractor networks. In B. Ross (Ed.). *Psychology of learning and motivation* (Vol. 45, pp. 41–86). Amsterdam: Elsevier.

McRae, K., Cree, G., Seidenberg, M., & 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., Westmacott, R., & de Sa, V. (1999). Further evidence for feature correlations in semantic memory. *Canadian Journal of Experimental Psychology, 53*, 360–373.

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.

Medin, D. L., & Smith, E. E. (1984). Concepts and concept formation. *Annual Review of Psychology, 35*, 113–138. doi:10.1146/annurev.ps.35.020184.000553.

Mirman, D., & Magnuson, J. (2009). The effect of frequency of shared features on judgments of semantic similarity. *Psychonomic Bulletin & Review, 16*, 671–677.

Moss, H. E., Rodd, J. M., Stamatakis, E. A., Bright, P., & Tyler, L. K. (2005). Anteromedial temporal cortex supports fine-grained differentiation among objects. *Cerebral Cortex, 15*, 616–627.

Moss, H., & Tyler, L. (1997). A category-specific impairment for non-living things in a case of progressive aphasia. *Brain and Language, 60*, 55–58.

Moss, H., Tyler, L., & Devlin, J. (2002). The emergence of category specific deficits in a distributed semantic system. In E. Forde & G. Humphreys (Eds.), *Category-specificity in brain and mind* (pp. 115–148). Sussex: Psychology Press.

Moss, H., Tyler, L., Durrant-Peatfield, M., & Bunn, E. (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., Tyler, L., & Jennings, F. (1997). When leopards lose their spots: Knowledge of visual properties in category-specific deficits for living things. *Cognitive Neuropsychology, 14*, 901–950.

Moss, H. E., Tyler, L. K., & Taylor, K. I. (2007). Conceptual structure. In G. Gaskell (Ed.), *Oxford handbook of psycholinguistics* (pp. 217–234). Oxford: Oxford University Press.

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

Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review, 92*(3), 289–316.

Murphy, G. L., & Wisniewski, E. J. (1989). Categorizing objects in isolation and in scenes: What a superordinate is good for. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15*(4), 572–586.

Pexman, P., Hargreaves, I., Siakaluk, P., Bodner, G., & Pope, J. (2008). There are many ways to be rich: Effects of three measures of semantic richness on visual word recognition. *Psychonomic Bulletin & Review, 15*, 161–167.

Pexman, P., Holyk, G., & Monfils, M. (2003). Number-of-features effects and semantic processing. *Memory and Cognition, 31*, 842–855.

Pexman, P., Lupker, S., & 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.

Pinheiro, J., & Bates, D. (2000). *Mixed effects models in S and S-plus*. New York: Springer.

Potter, M., & Faulconer, B. (1979). Understanding nonu phrases. *Journal of Verbal Learning and Verbal Behavior, 18*, 509–521.

Price, C., & Humphreys, G. (1989). The effects of surface detail on object categorization and naming. *The Quarterly Journal of Experimental Psychology: A, Human Experimental Psychology, 41*(4), 797–827.

R Development Core Team. (2008). *R: A language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing.

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.

Riddoch, M., & Humphreys, G. (1987). A case of integrative visual agnosia. *Brain, 110*, 1431–1462.

Rogers, T. T., Lambon Ralph, M. A., Garrard, P., Bozeat, S., McClelland, J. L., Hodges, J. R., et al. (2004). The structure and deterioration of semantic memory: A neuropsychological and computational investigation. *Psychological Review, 111*(1), 205–235.

Rosch, E., Mervis, C., Gray, W., Johnson, D., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology, 8*, 382–439.

Smith, E. E., & Medin, D. M. (1981). *Categories and concepts*. Cambridge, MA: Harvard University Press.

Snodgrass, J., & McCullough, B. (1986). The role of visual similarity in picture categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 12*, 147–154.

Stewart, F., Parkin, A., & Hunkin, N. (1992). Naming impairments following recovery from herpes simplex encephalitis: Category-specific? *Quarterly Journal of Experimental Psychology, 44A*, 261–284.

Tabossi, P. (1988). Effects of context on the immediate interpretation of unambiguous nouns. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 14*(1), 153–162. doi:10.1037/0278-7393.14.1.153.

Taylor, K. I., Moss, H. E., Randall, B., & Tyler, L. K. (2004). The interplay between distinctiveness and intercorrelation in the automatic activation of word meaning. In *Abstracts of the Psychonomic Society* (Vol. 9, p. 109).

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, UK: Cambridge University Press.

Taylor, K. I., Salamoura, A., Randall, B., Moss, H. E., & Tyler, L. K. (2008). Clarifying the nature of distinctiveness by domain interaction in conceptual structure: A comment on Cree, McNorgan and McRae (2006). *Journal of Experimental Psychology: Learning, Memory, and Cognition, 34*, 719–725.

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., Durrant-Peatfield, M., & Levy, J. (2000). Conceptual structure and the structure of concepts: A distributed account of category-specific deficits. *Brain and Language, 75*, 195–231.

Ulrich, R., & Miller, J. (1994). Effects of truncation on reaction-time analysis. *Journal of Experimental Psychology: General, 123*, 34–80.

Vigliocco, G., Vinson, D., Lewis, W., & Garrett, M. (2004). Representing the meanings of object and action words: The featural and unitary semantic space hypothesis. *Cognitive Psychology, 48*, 2004.

Vinson, D., Vigliocco, G., Cappa, S., & Siri, S. (2003). The breakdown of semantic knowledge: Insights from a statistical model of meaning representation. *Brain and Language, 86*, 347–365.

Warrington, E. (1975). Selective impairments of semantic memory. *Quarterly Journal of Experimental Psychology: General, 123*, 34–80.

Wittgenstein, L. (1953). *Philosophical investigations*. Oxford: Blackwell.