



Contents lists available at ScienceDirect

Journal of Memory and Language

journal homepage: www.elsevier.com/locate/jmlShared features dominate semantic richness effects for concrete concepts [☆]Ray Grondin, Stephen J. Lupker ^{*}, Ken McRae ^{*}

Department of Psychology, University of Western Ontario, London, Ont., Canada N6A 5C2

ARTICLE INFO

Article history:

Received 20 September 2006

Revision received 28 August 2008

Available online 19 October 2008

Keywords:

Semantic memory

Lexical concepts

Number-of-features effect

Modality-specific conceptual representations

ABSTRACT

When asked to list semantic features for concrete concepts, participants list many features for some concepts and few for others. Concepts with many semantic features are processed faster in lexical and semantic decision tasks [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; Pexman, P. M., Holyk, G. G., & MonFils, M.-H. (2003). Number-of-features effects and semantic processing. *Memory & Cognition*, 31, 842–855]. Using both lexical and concreteness decision tasks, we provided further insight into these number-of-features (NoF) effects. We began by replicating the effect using a larger and better controlled set of items. We then investigated the relationship between NoF and feature distinctiveness and found that features shared by numerous concrete concepts such as <has four legs> facilitate decisions to a greater extent than do distinctive features such as <moos>. Finally, we showed that NoF effects are carried by shared visual form and surface, encyclopedic, tactile, and taste knowledge. We propose a decision-making account of these results, rather than one based on the computation of word meaning.

© 2008 Elsevier Inc. All rights reserved.

People use language every day to convey messages, and inherent in our ability to understand these messages is our ability to compute the meaning of individual words. The computation of word meaning has been studied from a number of perspectives, including conceptual and linguistic development, the ability of normal adults to produce and understand language, neuropsychological impairments following brain injury, and computational modeling of all of these phenomena. One factor that has emerged as important in understanding the computation of word meaning is the richness of a word's semantic representation. Specifically, in many experimental tasks, participants respond more

quickly to words having richer semantic representations (Borowsky & Mason, 1996; de Groot, 1989; Hino & Lupker, 1996; James, 1975; Pexman, Lupker, & Hino, 2002). The goal of the present research is to further our understanding of the nature of word meaning by investigating how a particular aspect of semantic richness (the number of features a concept has) influences performance in two speeded tasks: lexical and concreteness decision. Study 1 is a replication of number-of-feature (NoF) effects using extremely carefully controlled materials. Studies 2 and 3 establish that feature distinctiveness matters in that the number of a concept's features that are shared by other concepts is an important determinant of NoF effects. Study 4 takes a modality-specific approach and shows that shared visual form and surface features, encyclopedic features, tactile features, and taste features underlie NoF effects. Finally, we argue that NoF effects are driven primarily by decision processes, rather than the activation of meaning per se.

Semantic richness

Semantic richness an important topic of study because understanding its impact has implications for theories of

[☆] This work was supported by a Natural Science and Engineering Research Council Master's Scholarship and a SHARCNET Graduate Research Fellowship to Ray Grondin, and National Science and Engineering Research Council Grant OGP0155704 and National Institute of Health Grant HD053136 to Ken McRae. We also thank Samantha Proffitt for her assistance with the experiments, Corey Rae McRae for her assistance with the feature type analyses, and Lawrence W. Barsalou for his comments.

^{*} Corresponding authors. Fax: +1 858 534 6788.

E-mail addresses: lupker@uwo.ca (S.J. Lupker), mcr@uwo.ca (K. McRae).

semantic memory. Some models of semantic memory assume a localist representation of concepts whereby the meaning of a word such as *robin* is represented by a single node in a network (Collins & Quillian, 1969; McNamara, 1992; Page, 2000). Other models, however, assume a distributed representation of concepts whereby a word's meaning is represented over multiple nodes in a network, each of which can be activated to varying degrees (Masson, 1995; Plaut & Shallice, 1993). Although the debate between local versus distributed representations is longstanding, with respect to semantic representations the balance of evidence appears to have tipped in favor of distributed models. (In fact, Page is careful to describe localist models as being characterized by the presence of localist representations, rather than the absence of distributed representations, thus blurring the distinction between the two types of models.) In general, the fact that words vary in terms of their degree of semantic richness, and the finding that semantic richness influences processing, seem more consistent with a distributed view of semantic memory because distributed representational systems more naturally encode graded representations. Furthermore, a distributed feature-based view of semantic memory leads to studying variables that would not be important in a localist scheme.

Semantic richness correlates with two central properties of word meaning, ambiguity and concreteness. Some words such as *bowl* are ambiguous in that they have multiple meanings, whereas others such as *tent* are unambiguous in that they have a single meaning. It has been assumed that ambiguous words have richer semantic representations because multiple meanings must be represented. Using lexical decision and naming tasks, Hino and Lupker (1996) found an ambiguity advantage whereby ambiguous words were responded to more quickly than unambiguous words (see also Jastrzembski, 1981; Kellas, Ferraro, & Simpson, 1988; Rubenstein, Garfield, & Millikan, 1970), suggesting that the meanings of ambiguous words may be computed more rapidly.

Interestingly, Hino, Pexman, and Lupker (2006) demonstrated that an ambiguity advantage does not arise in a semantic categorization task. Logically, however, the disappearance of the ambiguity advantage in these types of tasks should not be surprising given the decision-making problems created by ambiguous words. That is, the multiple meanings of a word typically do not support the same semantic decision (e.g., *bat*—is it an animal?), creating competition during the decision-making process. What this result does underline, therefore, is the importance of considering the impact of decision-making processes in many of these word recognition tasks, a point that is central to this article.

A second relevant finding is that response latencies for concrete words such as *robin* are shorter than for abstract words such as *justice* in tasks such as lexical decision (Binder, Westbury, McKiernan, Possing, & Medler, 2005) and naming (Strain, Patterson, & Seidenberg, 1995). A number of researchers have argued that this difference can be explained in terms of richer semantic representations for concrete words (Paivio, 1986; Plaut & Shallice, 1993).

An empirically based featural representation of word meaning

One way in which semantic richness effects has been studied is in terms of featural information. Although feature-based representations have proven fruitful in many ways, often their implementations are based on computer or experimenter-generated semantic representations and thus lack validation. To generate feature-based representations with increased psychological validity, a first approximation is to ask people what they know about the things to which various words refer. Of course, people cannot introspectively tell us everything that exists in their conceptual representations, but we can assume that what they do tell us provides a window into their actual underlying conceptual representations (Medin, 1989).

McRae, Cree, Seidenberg, and McNorgan (2005) adopted just such an approach. For both living (*robin*) and nonliving (*chair*) things, they had participants list descriptive features. For example, for *robin*, participants listed features such as <has wings>, <flies>, <eats worms>, and <has a red breast>. In total, 725 participants contributed to these semantic feature production norms for 541 concepts. This large set was used to define a psychological semantic space consisting of over 2000 features, and enabled the calculation of many statistics such as correlations between features and feature distinctiveness. In conjunction with various other measures (e.g., word frequency, word length, conceptual familiarity, orthographic and phonological neighborhoods), these empirically derived conceptual representations provide a rich basis for testing theories of semantic representation and computation.

In McRae et al.'s (2005) norms, each word refers to a concrete object (living or nonliving) and each has, as much as possible, a single meaning.¹ This minimizes potential differences among concepts, and allows a somewhat different definition and operationalization of semantic richness, the number of features (NoF) a concept is deemed to have. Importantly for this article, the norms also allow us to investigate the impact of the types of features that may drive these effects. That is, it may be the case that semantic richness effects are not simply due to additional features; rather, the types of features may matter. Furthermore, the types of features may interact with the manner in which the information is being used. In terms of cognitive experiments, the types of features may interact with the task under consideration.

¹ Although ambiguous concept names were avoided as much as possible, some do exist in the norms given that such a large proportion of English words are ambiguous. Noun–verb ambiguities such as *hammer* were dealt with by asking participants in the norming task to focus on the noun meaning, and the resulting features show that they did. Some noun–noun ambiguities were identified *a priori*, and a disambiguating cue was provided (e.g., *bat*—*baseball* and *bat*—*animal*). However, none of these items were used in the current studies. For ambiguities not identified by the researchers *a priori* (e.g., due to slang as in *pig*), participants rarely provided features relevant to the other meaning.

Probing NoF effects

Using McRae, de Sa, and Seidenberg's (1997) norms (which are a subset of McRae et al.'s, 2005, norms), Pexman et al. (2002) operationalized semantic richness as the number of features listed by participants. If richness underlies the processing advantage observed for ambiguous versus unambiguous words, and abstract versus concrete words, then the same advantage should be found when comparing words that differ in number of features.

Pexman et al. (2002) generated two sets of concepts using the 190 concepts in McRae et al.'s (1997) norms. One set contained 25 low number-of-features (low NoF) concepts, and the other contained 25 high number-of-features (high NoF) concepts. They also generated two sets of 50 nonword filler items. The first contained pronounceable nonwords whose spelling and sound do not correspond to any English word, such as *meap*, whereas the second contained pseudohomophones, nonwords whose spelling does not correspond to any English word, but whose sound does, such as *keap*. When the low and high NoF concepts were combined with the pronounceable nonwords, lexical decision latencies were shorter for high than for low NoF concepts. This effect was larger when pseudohomophones were used. Similar results were obtained when participants named the same low and high NoF items aloud.

Pexman, Holyk, and MonFils (2003) investigated the NoF effect using tasks more directly involving semantics. With the same stimuli described above, participants were faster to read high NoF words than low NoF words in a self-paced reading task when the sentence was incongruent ("After a heavy snowfall, Joel has to wear his *airplane*"), but not when it was highly congruent ("After the crash, Bob was nervous about getting on an *airplane*"), nor moderately congruent ("When I go home, I tend to travel by *airplane*"). In addition, using all 190 concrete concepts available at the time in McRae et al.'s (1997) norms, Pexman et al. had participants decide whether each word referred to a concrete object or to something abstract (e.g., *justice*). NoF was a significant predictor of decision latencies for concrete objects.

These findings invite questions regarding precisely what aspects of concepts are driving the NoF effects. That is, do all features matter equally? Do some types of features such as visual parts, functions, or characteristic behaviors, matter more than others? A number of recent studies demonstrate that not all types of features are created equal. For example, differential influences of whether a feature is distinctive to a concept, such as <moos> for cow, or is shared by many concepts, such as <has legs>, have been shown. Distinctiveness effects have been found in feature verification tasks (Cree, McNorgan, & McRae, 2006; Randall, Moss, Rodd, Greer, & Tyler, 2004), picture naming (Humphreys, Riddoch, & Quinlan, 1988), and semantic categorization of pictures (Riddoch & Humphreys, 1987). Given that distinctiveness matters in a number of tasks, it is possible that features influence NoF effects differentially depending on whether they are distinctive or shared. Furthermore, as is discussed in more detail in the introduction to Study 2, the influence of this factor var-

ies systematically as a function of the task. Therefore, in Studies 2 and 3, we tested whether the NoF effect depends on whether features are distinctive to a particular concept or are shared among concepts.

A second factor that has been shown to be important in numerous recent studies of semantic memory involves modality-specific aspects of concepts. That is, rather than concepts being represented in an amodal store in which information regarding the perceptual experience is abstracted away, recent evidence suggests that concepts are distributed across brain regions that represent various perceptual modalities. Functional magnetic resonance imaging studies demonstrate that modality-specific aspects of concrete concepts are activated when a picture is viewed or a word is read (Goldberg, Perfetti, & Schneider, 2006; Martin & Chao, 2001). Behavioral experiments have provided further evidence that concepts are distributed across modality-specific informational and neural regions (Pecher, Zeelenberg, & Barsalou, 2003). Modality specificity is also an important aspect of most accounts of category-specific semantic deficits (Garrard, Lambon Ralph, Hodges, & Patterson, 2001; Simmons & Barsalou, 2003).

Using linguistically based features for establishing factors such as the NoF seems, at first blush, to correspond to assuming amodal representations. However, features can be classified as belonging to specific perceptual modalities, and thus be used as a basis for modality-specific representations. For example, Cree and McRae (2003) classified all of the features in McRae et al.'s (2005) norms into a nine-way feature type taxonomy based on the current state of the art in neuroscience and cognitive neuropsychology. These feature types included, for example, visual form and surface features, tactile features, functional features, taste features, and encyclopedic features. Their relative salience was crucial in accounting for the behavioral trends observed in patients with category-specific semantic deficits. Given the centrality of modality-specific aspects of concepts to conceptual processing, it seems likely that features indexing different types of knowledge may play differential roles in the facilitation found in studies of NoF effects. For example, people's knowledge about how a living or nonliving thing looks might be particularly important information, and thus play a critical role. Therefore, in Study 4, we analyzed whether features that correspond to different types of knowledge contribute differentially to NoF effects. These analyses provide insight into both NoF effects and the nature of semantic representations.

Study 1

The purpose of Study 1 was to establish further the baseline empirical phenomena by testing whether Pexman et al.'s (2002, 2003) results replicate in both lexical and concreteness decision tasks using improved sets of items. A close investigation of the items used in Pexman et al.'s (2002, 2003) studies reveals that some variables known to influence word processing (word frequency and word length) were not perfectly equated between the high and low NoF word sets, and that the differences favored the

high NoF items. Although Pexman et al. addressed this issue by partialing out the influence of these variables using multiple regression, it may be that the observed NoF effects were due to the combined influence of these confounded variables. It was possible to construct larger lists of concepts (64 versus 25) that are better balanced on a larger set of variables because there now exists a larger set of norms (McRae et al., 2005, containing 541 concepts versus McRae et al., 1997, containing 190 concepts).

Study 1A

Methods

Participants. Seventeen undergraduate students at the University of Western Ontario received partial course credit for their participation. In all experiments reported herein, participants had either normal or corrected-to-normal visual acuity and were native English speakers.

Materials. Two sets of words referring to concrete objects were generated from McRae et al.'s (2005) norms, 64 low and 64 high NoF concepts (see Appendix A). High NoF concepts contained a significantly greater number of features than did low NoF concepts, $t(126) = 29.67, p < .001$. NoF can be calculated by either including or excluding taxonomic features (e.g., <is an animal>, <is a vegetable>), which seem qualitatively different than the vast majority of other features, such as parts, colors, functions, characteristic behaviors, and so on. Therefore, we excluded taxonomic features from the counts in the present analyses. Note that in Pexman et al.'s (2002, 2003) studies, NoF was calculated including taxonomic features, but it was later found that NoF without taxonomic features had a stronger relationship to decision latencies and decision errors in a concreteness decision task (Pexman & Hope, in preparation).

The two sets of items were matched extremely closely on variables known to influence lexical and concreteness decisions (see Table 1). These included word frequency, which was computed using the natural logarithm of the singular plus plural counts taken from the British National Corpus (BNC) online search engine (Burnard, 2000).

Table 1
Characteristics of stimuli used in Studies 1A and 1B

Variable	Low NoF		High NoF	
	<i>M</i>	<i>SE</i>	<i>M</i>	<i>SE</i>
Number of features (NoF)	9.0	0.2	15.7	0.3
ln(BNC) frequency	6.4	0.2	6.4	0.2
Familiarity	5.7	0.2	5.7	0.2
Number of letters	5.4	0.2	5.4	0.2
Number of phonemes	4.4	0.2	4.5	0.2
Number of syllables	1.6	0.1	1.7	0.1
Orthographic neighborhood size (N)	4.6	0.8	4.4	0.7
Semantic density	157	20	155	11
Number of creatures	18		21	
Number of fruits/vegetables	11		8	
Number of living things	35		35	
Number of musical instruments	0		0	

Note. ln = the natural logarithm (\log_e); BNC = British National Corpus.

Although Pexman and colleagues computed their frequencies using the Kučera and Francis (1967) corpus, we chose the BNC because it is based on a larger corpus (89.7 versus 1 million words) and is more recent (2000 versus 1967). Concept familiarity was measured by asking 20 participants to rate "How familiar are you with the thing that the word refers to?" on a 9-point scale, with 1 corresponding to *not at all familiar*, and 9 corresponding to *highly familiar*. Number of letters, number of phonemes, number of syllables, and orthographic neighborhood size (Coltheart, Davelaar, Jonasson, & Besner, 1977) were computed using the N-Watch program (Davis, 2005).

The concepts were also equated on semantic density as calculated from McRae et al.'s (2005) norms. In the norms, each feature is a vector of production frequencies (the number of participants listing that feature for each specific concept) across the 541 concepts. Proportion of shared variance for each pair of features was calculated by squaring the correlation between the vectors for the two features. Only features occurring in three or more concepts were included to attempt to avoid spurious correlations. A concept's semantic density was calculated as the sum of the proportion of shared variances for each pair of features that are included in that concept. Thus, semantic density provides a measure of the degree to which a concept's features are intercorrelated. Finally, because there is evidence that different categories of concrete objects are processed differentially (Laws & Gale, 2002), we also matched the two groups according to the following category breakdown: creatures, fruits and vegetables, nonliving things, and musical instruments.

Filler items consisted of 128 pronounceable nonwords whose spelling and sound did not correspond to an English word (*lerve*). The nonwords and low and high NoF concept names were matched on mean number of letters so that length did not cue the response.

Procedure. Participants were tested individually using PsyScope (Cohen, MacWhinney, Flatt, & Provost, 1993) on a Macintosh computer equipped with a 16-in. monitor and a CMU button box (which provides decision latencies with accuracy to the nearest ms). Letters were approximately 0.5 cm high, black, and presented on a white background. One item was presented at a time and participants were asked to decide whether the presented item referred to an English word or not. They used the index finger of their dominant hand for a 'yes' response and the index finger of their non-dominant hand for a 'no' response. Decision latencies were measured from stimulus onset to the onset of the button press. Participants were instructed to respond as quickly and accurately as possible.

Participants first completed 30 practice trials with verbal feedback concerning incorrect decisions. Items presented during the practice trials did not appear in the test trials. On each trial, the item was presented until the participant made a decision. The intertrial interval was 1500 ms. Test trials were identical to the practice trials except that no feedback was provided and there was a break after every 50 items. Items were presented in random order. The experiment took approximately 25 min.

Results and discussion

Subject (t_1) and item (t_2) analyses were performed on decision latencies and the square root of the number of errors (Myers, 1979). Errors were removed from decision latency analyses (4.6% of trials). Correct decisions exceeding three standard deviations above the grand mean for the target words were replaced with the cutoff value (1.8% of trials). NoF (high versus low) was within-subjects (paired t -test) and between-items (independent t -test). All p -values assume a two-tailed distribution. Mean decision latencies and error rates are presented in Table 2.

Lexical decision latencies to high NoF concepts were 30 ms shorter than to low NoF concepts, $t_1(16) = 7.37$, $p < .00001$, $t_2(126) = 2.34$, $p < .05$, $\min F(1, 141) = 4.97$, $p < .05$. Participants made 3.7% fewer errors on high NoF concepts, $t_1(16) = 5.99$, $p < .0001$, $t_2(126) = 2.03$, $p < .05$, $\min F(1, 142) = 3.69$, $p < .06$. For nonwords, the mean decision latency was 756 ms, and the mean error rate was 3.7%. In summary, using this improved set of items, we replicated Pexman et al.'s (2002) NoF effect in the lexical decision task.

Study 1B

Methods

Participants. Seventeen undergraduate students at the University of Western Ontario received partial course credit for their participation.

Materials. Target items were identical to those in Study 1A. Filler items were 128 abstract concepts matched with the target items on mean number of letters. Although some abstract filler items were noun/verb ambiguous (*respect*), they were chosen to have a salient noun meaning.

Procedure. The procedure was identical to that in Study 1A except that for each presented word, participants were asked to respond 'yes' if the thing to which the word referred was touchable (*robin*) and 'no' otherwise (*justice*).

Results and discussion

The analyses were identical to those in Study 1A. Errors were removed from the decision latency analyses (3.9% of

trials). Correct decisions that exceeded three standard deviations above the concrete concept grand mean were replaced with the cutoff value (1.4% of trials). Mean decision latencies and error rates are presented in Table 2.

Concreteness decision latencies to high NoF concepts were 29 ms shorter than to low NoF concepts, $t_1(16) = 4.27$, $p < .0001$, $t_2(126) = 2.51$, $p < .05$, $\min F(1, 118) = 4.68$, $p < .05$. Participants also made 3.5% fewer errors on high NoF concepts, which was significant by subjects and marginal by items, $t_1(16) = 5.84$, $p < .0001$, $t_2(126) = 1.80$, $p < .08$, $\min F(1, 141) = 2.95$, $p < .09$. The mean decision latency for abstract concepts was 706 ms, and the mean error rate was 4.7%. Thus, Study 1B replicated the NoF effect using a concreteness decision task with an extremely tightly controlled and large set of items.

In conjunction with Pexman et al.'s (2002, 2003) results, the present results provide strong evidence supporting the claim that decisions in speeded tasks are faster for concepts with a greater number of semantic features. In Studies 2 and 3, we examined the NoF effect more closely. Given previous results, described more fully below, it is quite possible that the NoF effect is mediated by feature type. Specifically, the possibility exists that features that are shared by many concepts (<made of metal>) versus those which are highly distinctive to a concept (<moos>) may be differentially responsible for NoF effects.

Shared versus distinctive features

In their investigation of the factors underlying trends in the performance of category-specific deficit patients, Cree and McRae (2003) included a measure of feature distinctiveness (see also Garrard et al., 2001). At one end of the distinctiveness dimension lie highly *shared* features such as <has four legs> and <is hard> that occur in many concepts. On the other end lie *distinctive* features such as <moos> and <oinks> that occur in few concepts (or even just one). Shared features denote commonalities among concepts, and thus indicate similarities among concepts. In contrast, distinctive features denote differences, and thus help people to discriminate among concepts. Cree and McRae defined a feature as shared if it was listed for more than 2 of the 541 concepts, and distinctive if it was listed for only 1 or 2 concepts. They also computed feature distinctiveness as a continuous dimension (1/number of concepts in which a feature occurred), but for explanatory purposes, we focus on the shared versus distinctive binary measure.

Cree and McRae (2003) showed that this measure of distinctiveness was a key part of understanding behavioral trends across category-specific deficit patients. These deficits are typically observed in tasks such as picture naming, word-to-picture matching, defining, and naming from definition, all of which require distinguishing a concept from among similar concepts. For example, a patient can only identify something as a zebra rather than a horse in a picture naming task if their knowledge about what distinguishes zebras from horses (e.g., black and white stripes) is preserved. Cree and McRae computed the proportion of distinctive features for a large number of concepts. Assuming that damage to features (distinctive versus shared) is equiprobable, the implication is that concepts with lower

Table 2

Decision latencies (ms) and error rates (%) for Studies 1A (lexical decision) and 1B (concreteness decision)

	Study 1A		Study 1B	
	<i>M</i>	<i>SE</i>	<i>M</i>	<i>SE</i>
<i>Latencies</i>				
Low NoF	623	19	666	16
High NoF	593	20	637	14
Difference	30**		29**	
<i>Error rates</i>				
Low NoF	6.5	0.8	5.6	0.9
High NoF	2.8	0.5	2.1	0.6
Difference	3.7**		3.5*	

Note. NoF = number of features.

* Significant by subjects.

** Significant by subjects and items.

proportions of distinctive features should be impaired to a greater extent than those with higher proportions in these types of tasks. Indeed, the degree to which concepts' features are distinctive accounted for the degree of impairment across multiple categories.

Analogous results have been found with normal adults. Humphreys et al. (1988) found that participants were faster to name pictures of objects belonging to categories whose exemplars are structurally dissimilar (clothing and furniture) versus structurally similar (insects, fruits, and vegetables). In feature verification tasks, Cree et al. (2006) found that participants were faster to verify a concept-feature pair when the feature was distinctive than when it was shared. They simulated these results using an attractor network and found that distinctive features had stronger connections to and from other features of the same concept than did shared features, which provides a plausible basis for the human results.

There is also, however, research showing what superficially appears to be the opposite pattern, an advantage for items that share many features with other items. For example, typical exemplars, such as *robin* for the category bird, which possess a greater number of shared features than do atypical exemplars such as *ostrich*, are responded to more rapidly in a category verification task (Smith, Shoben, & Rips, 1974). Similarly, in contrast to their findings using picture naming, Riddoch and Humphreys (1987) found that participants were faster to make broad-level classifications (living versus nonliving) for pictures of objects belonging to categories whose exemplars are structurally similar.

In resolving this apparent contradiction, what is important to note is that broad semantic classification tasks do not require people to make distinctions among category exemplars, and it is only in tasks in which participants are not required to make such distinctions that concepts with large numbers of shared features show a processing advantage. In essence then, the processing advantages for distinctive versus shared feature types are task dependent. Distinctive features make it easier to respond when the task requires distinguishing an item from among similar items, such as when naming the picture of an object. Shared features, particularly those consistent with the decision required, make it easier to respond when a more broad-based judgment is called for, such as deciding whether a *robin* is a living thing.

Concreteness decision would appear to be a task that calls for a more broad-based judgment. Thus, one might expect that shared features would be more important than distinctive features in that task. Whether the same is true in lexical decision depends, first of all, on the extent to which semantic information is used. If semantic information does play a significant role in lexical decision, then the type of semantic information that is activated may matter. In particular, if the sets of activated features are consistent with a large number of concepts, that fact may provide good evidence that the letter string corresponds to a real word. As a result, shared features may contribute more than distinctive features to the NoF effect in the lexical decision task as well.

In our Study 1 stimuli, low NoF concepts averaged 9.0 features, whereas high NoF concepts averaged 15.7 features. Breaking these totals down, low NoF concepts averaged 2.1 distinctive and 6.9 shared features, whereas high NoF concepts averaged 5.9 distinctive and 9.8 shared features. In Pexman et al.'s (2002) studies, low NoF concepts averaged 11.0 features, whereas high NoF concepts averaged 18.1 features. Low NoF concepts averaged 3.5 distinctive and 7.5 shared features, whereas high NoF concepts averaged 5.9 distinctive and 12.2 shared features. Because the high and low NoF concepts differed in terms of both distinctive and shared features in both cases, Study 1 and Pexman et al.'s experiments do not provide insight into their relative contributions.

One way to investigate the relative contributions of shared versus distinctive features is to directly contrast the number of each type of feature while holding NoF constant. Another way is to create high and low NoF conditions by independently manipulating the number of shared or the number of distinctive features while holding the other type constant. Study 2 used the former manipulation whereas Study 3 used the latter.

Study 2

The purpose of Study 2 was to investigate whether lexical and concreteness decisions are systematically influenced when the number of shared versus distinctive features is varied while holding NoF constant.

Study 2A

Methods

Participants. Twenty-five undergraduate students at the University of Western Ontario received \$10 for their participation.

Materials. Two sets of target words referring to concrete objects were generated from McRae et al.'s (2005) norms. One set consisted of 55 low number of shared features concepts and the other consisted of 55 high number of shared features concepts (see Appendix B). High shared-feature concepts had a significantly greater number of shared features, $t(108) = 12.07, p < .001$. The two sets were matched on the same variables described in Study 1, plus NoF (see Table 3). Filler items consisted of 110 nonwords, which were a subset of those used in Study 1A and were matched with the target items on number of letters.

Procedure. A lexical decision task identical to that in Study 1A was used.

Results and discussion

Subject (t_1) and item (t_2) analyses were performed on decision latencies and error rates. The independent variable was number of shared features, which was within-subjects but between-items. Errors were removed from the decision latency analyses (3.6% of trials). Correct decisions exceeding three standard deviations above the grand mean of the target words were replaced with the cutoff

Table 3
Characteristics of stimuli used in Studies 2A and 2B

Variable	Number of shared features			
	Low		High	
	M	SE	M	SE
Number of shared features	6.7	0.2	10.7	0.2
Number of distinctive features	6.0	0.3	1.7	0.1
Number of features (NoF)	12.7	0.4	12.4	0.3
ln(BNC) frequency	6.6	0.2	6.5	0.2
Familiarity	6.0	0.3	6.0	0.3
Number of letters	5.4	0.2	5.4	0.2
Number of phonemes	4.5	0.2	4.6	0.2
Number of syllables	1.6	0.1	1.7	0.1
Orthographic neighborhood size (N)	4.4	0.7	4.3	0.8
Semantic density	170	19	172	13
Number of creatures	12		15	
Number of fruits/vegetables	5		6	
Number of living things	37		32	
Number of musical instruments	1		2	

Note. ln = the natural logarithm (\log_e); BNC = British National Corpus.

value (1.6% of trials). Mean decision latencies and error rates are presented in Table 4.

Lexical decision latencies for concepts with a high number of shared features were 11 ms shorter than for those with a low number of shared features, which was significant by subjects, $t_1(24) = 4.60$, $p < .001$, but not by items, $t_2(108) = 1.07$, $p > .2$, $\min F(1, 118) = 1.09$, $p > .2$. The 0.7% difference in error rates was nonsignificant, $t_1(24) = 1.32$, $p > .1$, $t_2(108) = 0.30$, $p > .7$, $\min F(1, 118) = 0.09$, $p > .7$. The mean decision latency for nonwords was 633 ms, and the mean error rate was 2.7%.

In summary, although the effect in the lexical decision task was not large, there is evidence that concepts with numerous shared features are easier to respond to than those with fewer shared features. This effect may be significant only in the subject analyses for two reasons. First, the feature type manipulation was a within-subject but between-item manipulation. Within manipulations are inevitably more powerful than between manipulations. Second, as alluded to earlier, the lexical decision task is not based solely on semantic knowledge. Instead, lexical decisions are made on the basis of some combination of orthographic, phonological, and semantic knowledge, and

Table 4
Decision latencies (ms) and error rates (%) for Studies 2A (lexical decision) and 2B (concreteness decision)

	Study 2A		Study 2B	
	M	SE	M	SE
<i>Latencies</i>				
Low shared	555	13	756	31
High shared	544	12	714	29
Difference	11*		42**	
<i>Error rates</i>				
Low shared	3.9	0.6	4.5	0.8
High shared	3.2	0.5	3.0	0.6
Difference	0.7		1.5	

* Significant by subjects.

** Significant by subjects and items.

the influence of these types of information can be modulated by task parameters (such as the type of nonword foils used). Unless semantic knowledge plays a major role, one would not expect the impact of a semantic factor like distinctiveness to be particularly strong. On the other hand, a semantic decision task such as the concreteness decision task used in Study 1B unambiguously depends on the computation of word meaning, which should make any NoF effect easier to detect. Therefore, Study 2B used the same stimuli as Study 2A in a concreteness decision task.

Study 2B

Methods

Participants. Twenty-four undergraduate students at the University of Western Ontario received \$10 for their participation.

Materials. The test items were identical to those used in Study 2A. Filler items consisted of 110 abstract concepts. The fillers were a subset of those used in Study 1B and were matched with the target items on number of letters.

Procedure. A concreteness decision task identical to that in Study 1B was used.

Results and discussion

The analyses were identical to those in Study 2A. Errors were removed from the decision latency analyses (3.8% of trials). Correct decisions that exceeded three standard deviations above the concrete word grand mean were replaced with the cutoff value (1.6% of trials). Mean decision latencies and error rates are presented in Table 4.

Concreteness decision latencies for concepts with a high number of shared features were 42 ms shorter than for those with a low number of shared features, $t_1(23) = 6.17$, $p < .00001$, $t_2(108) = 2.18$, $p < .05$, $\min F(1, 127) = 4.22$, $p < .05$. The 1.5% difference in error rates was nonsignificant, $t_1(23) = 1.10$, $p > .2$, $t_2(108) = 1.16$, $p > .2$, $\min F(1, 127) = 0.63$, $p > .6$. The mean decision latency for abstract concepts was 791 ms, and the mean error rate was 4.8%. Thus, increasing the number of shared features leads to faster concreteness decisions, with the size of the effect being greater than in Study 2A with lexical decision.

Obtaining a larger effect in semantic decision is not particularly surprising. Although there was no difference in effect size between Study 1A with lexical decision (30 ms) and Study 1B with concreteness decision (29 ms), a number of studies have found stronger effects of semantic manipulations in semantic rather than lexical decision tasks (Becker, Moscovitch, Behrmann, & Joordens, 1997; Bueno & Frenck-Mestre, 2008). It is clear that participants must retrieve semantic information in order to decide whether a word refers to something that is concrete or abstract. In contrast, computation of word meaning is, undoubtedly, less strongly related to making a word versus nonword decision (Pexman et al., 2002) because orthographic and phonological information undoubtedly also play a role. Thus, one would expect semantic effects to be smaller in lexical than concreteness decision tasks.

One question that Study 2 does not allow us to answer is, assuming that the impact of additional shared features is a processing benefit, what is the impact of additional distinctive features? If the issue is merely semantic richness, one would expect that adding features, regardless of type, would speed processing because, in all cases, the result is a richer semantic representation. However, it is possible that adding distinctive features inhibits certain decisions because as distinctive features are added to a concept, that concept becomes increasingly less similar to other concepts. For example, such might be the case in a task such as semantic verification in which latencies are longest for concepts that are dissimilar to other category members (Rips, Shoben, & Smith, 1973). In fact, although we have discussed the manipulation in Study 2 as being a manipulation of shared features, if we reverse the logic and think about these data in terms of distinctive features, one could argue that distinctive features inhibit decisions. That is, because each feature was classified as either shared or distinctive, the numbers of shared and distinctive features are the complements of one another in Study 2. Re-labeling the low-shared condition as high-distinctive and the high-shared condition as low-distinctive suggests that decision latencies are longer to concepts with many versus few distinctive features. Although this interpretation is possible, it is unlikely given that in Study 1 there were, on average, a greater number of additional distinctive features (3.9) than additional shared features (2.9) in the high NoF condition and an NoF effect was obtained. Study 3 was designed to overcome this ambiguity of interpretation.

Study 3

The purpose of Study 3 was to investigate whether lexical and concreteness decisions are systematically influenced when the number of shared features and the number of distinctive features are manipulated while keeping the other constant. Therefore, shared and distinctive features were independently manipulated creating two separate manipulations of NoF, one based solely on

adding shared features, and the other based solely on adding distinctive features.

Study 3A

Methods

Participants. Forty-nine undergraduate students at the University of Western Ontario received partial course credit for their participation. Four participants were excluded for making more than 15% incorrect decisions across all items, leaving 45 participants.

Materials. Four sets of 20 concepts were generated from McRae et al.'s (2005) norms (see Table 5). In the first two sets, the number of shared features was manipulated while holding constant the number-of-distinctive-features. In the other two sets, the number-of-distinctive-features was manipulated while holding constant the number of shared features. High NoF, shared-manipulated concepts had a significantly greater number of shared features when compared to low NoF, shared-manipulated concepts, $t(38) = 18.45, p < .001$. High NoF, distinctive-manipulated concepts had a significantly greater number of distinctive features when compared to low NoF, distinctive-manipulated concepts, $t(38) = 16.83, p < .001$. The four sets were matched on the same variables described in Study 1. However, note that because semantic density, which measures the degree to which a concept's features are intercorrelated, is calculated using only shared features, it was lower for the low NoF, shared-manipulated condition that contains only 3.1 shared features on average. Filler items consisted of 80 nonwords taken from those used in Study 1 and were matched with the target items on number of letters.

Note that the two sets of conditions mirror one another. In the shared features manipulated conditions, there were 6.5 and 6.1 distinctive features on average, with the number of shared features being low (3.1 in the low NoF items) or high (9.1 in the high NoF items). In the distinctive features manipulated conditions, there were 6.5 and 6.3

Table 5
Characteristics of stimuli used in Studies 3A and 3B

Variable	Shared features manipulated				Distinctive features manipulated			
	Low NoF		High NoF		Low NoF		High NoF	
	<i>M</i>	<i>SE</i>	<i>M</i>	<i>SE</i>	<i>M</i>	<i>SE</i>	<i>M</i>	<i>SE</i>
Number of shared features	3.1	0.2	9.1	0.2	6.5	0.3	6.3	0.4
Number of distinctive features	6.5	0.6	6.1	0.5	2.9	0.2	9.1	0.3
Number of features (NoF)	9.5	0.6	15.2	0.6	9.4	0.3	15.4	0.3
ln(BNC) frequency	7.0	0.4	6.9	0.3	6.8	0.3	6.6	0.4
Familiarity	6.3	0.5	6.2	0.4	6.2	0.4	6.4	0.4
Number of letters	5.8	0.5	5.7	0.3	5.8	0.3	5.9	0.4
Number of phonemes	4.7	0.4	4.9	0.2	4.5	0.3	5.0	0.3
Number of syllables	1.9	0.2	1.6	0.1	1.6	0.1	1.9	0.1
Orthographic neighborhood size (N)	4.3	1.3	3.1	0.8	4.1	1.1	3.2	0.9
Semantic density	12	6	165	29	114	30	83	17
Number of creatures	0		4		4		1	
Number of fruits/vegetables	4		2		3		1	
Number of living things	15		14		13		18	
Number of musical instruments	1		0		0		0	

Note. ln = the natural logarithm (\log_e); BNC = British National Corpus.

shared features on average, with the number of distinctive features being low (2.9 in the low NoF items) or high (9.1 in the high NoF items).

Procedure. A lexical decision task identical to those in Studies 1A and 2A was used.

Results and discussion

Subject (F_1) and item (F_2) analyses of variance were conducted on decision latencies and error rates. The independent variables were type of manipulated feature (shared versus distinctive) and NoF (high versus low), both of which were within-subjects, but between-item variables. Planned contrasts were used to test for differences among all pairs of conditions. Errors were removed from the decision latency analyses (3.3% of trials). Correct decisions that exceeded three standard deviations above the grand mean of the target words were replaced with the cutoff value (1.7%). Mean decision latencies and percent errors are presented in Table 6.

Decision latencies. Manipulated feature type interacted with NoF by subjects, $F_1(1,44) = 7.59$, $p < .01$, $F_2(1,76) = 1.27$, $p > .2$, $\text{min}F(1,99) = 1.08$, $p > .2$. For the planned comparisons, the halfwidth of the confidence interval of the difference between means was 9 ms. If the halfwidth is less than the observed difference between means, then the contrast would be significant by a conventional inferential test. Confidence intervals were computed using the methods recommended by Masson and Loftus (2003), using pooled error terms because the error terms were similar in magnitude. Decision latencies in the low NoF, shared-manipulated condition were a significant 25 ms longer than in the high NoF, shared-manipulated condition. However, there was no difference when the number of distinctive features were manipulated (i.e., when the number of shared features was held constant). In fact, decision latencies in the low NoF, shared-manipulated condition, which had only 3.1 shared features on average, were significantly longer than in the other three conditions. There were no significant differences among those other three conditions, which differed by a maximum of 4 ms.

Collapsed across manipulated feature type, decision latencies were 10 ms shorter for high ($M = 598$ ms, $SE = 9$ ms) than for low NoF concepts ($M = 608$ ms,

$SE = 9$ ms), which was significant by subjects, $F_1(1,44) = 5.65$, $p < .05$, but not by items, $F_2 < 1$, $\text{min}F < 1$. Collapsed across NoF, decision latencies were 10 ms shorter for concepts with distinctive features manipulated ($M = 598$ ms, $SE = 8$ ms) than for those with shared features manipulated ($M = 608$ ms, $SE = 9$ ms), which was significant by subjects, $F_1(1,44) = 6.37$, $p < .05$, but not by items, $F_2 < 1$, $\text{min}F < 1$. Again, this effect was carried by the condition containing only 3.1 shared features.

Error rates. The interaction between manipulated feature type and NoF was not significant, $F_1(1,44) = 2.26$, $p > .1$, $F_2 < 1$, $\text{min}F < 1$. Collapsed across feature type, there was no significant difference in error rates between high ($M = 2.8\%$, $SE = 0.4\%$) and low NoF concepts ($M = 3.8\%$, $SE = 0.6\%$), $F_1 < 1$, $F_2 < 1$, $\text{min}F < 1$. Collapsed across NoF, there was no significant difference in error rates for distinctive ($M = 2.8\%$, $SE = 0.4\%$) versus shared features manipulated conditions ($M = 3.6\%$, $SE = 0.5\%$), $F_1(1,44) = 1.03$, $p > .3$, $F_2 < 1$, $\text{min}F < 1$. The mean decision latency for non-words was 756 ms, and the mean error rate was 6.4%. We defer discussion of these results until the discussion of Study 3B.

Study 3B

Methods

Participants. Forty-seven undergraduate students at the University of Western Ontario received partial course credit for their participation. One participant was excluded for talking during the study, and two were excluded for making greater than 15% errors, leaving 44 participants.

Materials. The test items were identical to those in Study 3A. Filler items were 80 abstract concepts that were a subset of those used in Study 1B and were matched with the target items on number of letters.

Procedure. A concreteness decision task identical to those in Studies 1B and 2B was used.

Results and discussion

The design and analyses were identical to Study 3A. Errors were removed from the decision latency analyses (4.7% of trials). Decision latencies that exceeded three standard deviations above the concrete word grand mean were replaced with the cutoff value (1.7%). Mean decision latencies and percent errors are presented in Table 6.

Decision latencies. The interaction between manipulated feature type and NoF was significant by subjects, $F_1(1,43) = 15.98$, $p < .001$, $F_2(1,76) = 2.54$, $p > .1$, $\text{min}F(1,98) = 2.19$, $p > .1$. The halfwidth of the confidence interval of the difference between means was 13 ms. Decision latencies in the low NoF, shared-manipulated condition were a significant 81 ms longer than in the high NoF, shared-manipulated condition. In fact, decision latencies in the low NoF, shared-manipulated condition (with only 3.1 shared features on average) were longer than in each of the other two conditions as well. When the number of shared features was held constant, and distinctive features

Table 6
Decision latencies (ms) and error rates (%) for Studies 3A and 3B

Dependant variable	Shared features manipulated				Distinctive features manipulated			
	Low NoF		High NoF		Low NoF		High NoF	
	M	SE	M	SE	M	SE	M	SE
<i>Study 3A (lexical decision)</i>								
Latencies	621	13	596	14	596	12	600	12
Error rates	4.8	0.9	2.4	0.5	2.8	0.6	3.2	0.5
<i>Study 3B (concreteness decision)</i>								
Latencies	792	19	711	18	736	19	714	19
Error rates	7.5	1.2	3.0	0.6	4.4	1.1	3.9	0.9

Note. NoF = number of features.

were manipulated, there was a significant 22 ms difference. The low NoF, distinctive-manipulated condition was also a significant 25 ms longer than the high NoF, shared-manipulated condition. Finally, there was only a 3 ms difference between the two high NoF conditions.

Collapsed across feature type, decision latencies were 51 ms shorter for high ($M = 713$ ms, $SE = 13$ ms) than for low NoF concepts ($M = 764$ ms, $SE = 14$ ms), $F_1(1,43) = 53.59$, $p < .00001$, $F_2(1,76) = 7.67$, $p < .01$, $\min F(1,96) = 6.71$, $p < .05$. Collapsed across NoF, decision latencies were 27 ms shorter for concepts with distinctive features manipulated ($M = 725$ ms, $SE = 13$ ms) than for those with shared features manipulated ($M = 752$ ms, $SE = 14$ ms), which was significant by subjects, $F_1(1,43) = 26.16$, $p < .00001$, but not by items, $F_2(1,76) = 2.12$, $p > .1$, $\min F(1,88) = 1.96$, $p > .1$.

Error rates. The interaction between manipulated feature type and NoF was significant by subjects, $F_1(1,43) = 6.43$, $p < .05$, $F_2(1,76) = 2.17$, $p > .1$, $\min F(1,113) = 1.62$, $p > .2$. Collapsed across feature type, participants made 2.6% fewer errors to high ($M = 3.4\%$, $SE = 0.8\%$) than to low NoF concepts ($M = 6.0\%$, $SE = 0.5\%$), $F_1(1,43) = 17.95$, $p < .0001$, $F_2(1,76) = 4.98$, $p < .05$, $\min F(1,109) = 3.90$, $p < .06$. Collapsed across NoF, there was not a significant difference when distinctive ($M = 4.1\%$, $SE = 0.7\%$) versus shared features were manipulated ($M = 5.2\%$, $SE = 0.7\%$), $F_1(1,43) = 2.92$, $p > .09$, $F_2 < 1$, $\min F < 1$. The mean decision latency for abstract words was 798 ms, and the mean error rate was 4.4%.

Study 3 shows that shared features play a major role, whereas distinctive features play a facilitatory but lesser role, particularly when the effects of NoF, shared features, and distinctive features are contrasted. The main effect of NoF was 10 ms in lexical decision and 51 ms in concreteness decision. The influence of distinctive features when shared features were equated was basically nil in lexical decision, but 22 ms in concreteness decision. The largest effects were found for shared features when the number of distinctive features was held constant, 25 ms in lexical decision and 81 ms in concreteness decision.

In both Study 3A and 3B, when NoF is low (9.5 features) and there are very few shared features (3.1), decision latency is consistently the longest. When NoF is high, decision latency is consistently the shortest. Note that the difference in the number of shared features for the two high NoF groups in Study 3 was smaller than in Study 2, and the results differed between experiments. In Study 2, the low-shared and high-shared groups contained 12.7 and 12.4 features, respectively, and the difference in number of shared features was 4.0 (6.7 versus 10.7, so that the high-shared items contained only 1.7 distinctive features on average). In Study 2, there was an effect of this difference in number of shared features, particularly in concreteness decision. In contrast, in Study 3, the high NoF, shared-manipulated and high NoF, distinctive-manipulated item sets contained 15.2 and 15.4 features, differing in number of shared features by only 2.8 (9.1 versus 6.3). No difference was obtained in either task in Study 3.

The other condition in Study 3 included the low NoF, distinctive-manipulated concepts. These items averaged 6.5

shared and only 2.9 distinctive features. In the concreteness decision latencies, this group patterned the way one would imagine. That is, it fell between the low NoF, shared-manipulated condition on the one hand (which contained 3.4 fewer shared features), and the two high NoF conditions on the other hand, which contained at least as many shared features, and a greater number of distinctive features. In other words, these results illustrate the influence of shared features and suggest that there is a positive impact of distinctive features as well. The one perplexing data point in Study 3 is the mean lexical decision latency for the low NoF, distinctive-manipulated condition. In this case, it was virtually identical to the two high NoF groups of items, showing a lack of influence of overall NoF. Note however that lexical decision latency for this group was shorter than for the low NoF, shared-manipulated items, thus showing an influence of the number of shared features.

Overall regressions. The above analyses support our conclusion regarding the importance of shared features to the NoF effect while suggesting that distinctive features also play a positive role. To provide a further examination of the relative influences of shared and distinctive NoF in the present experiments, we conducted stepwise regression analyses to compare their ability to predict decision latencies when the items from Studies 1–3 were combined. For the items that appeared in multiple experiments, we calculated the mean latency across experiments. We forced in $\ln(\text{BNC})$ word frequency and word length in letters on the first step to account for basic word-reading non-semantic variables. For concreteness decision, shared NoF significantly predicted decision latency, partial $r = -.26$, $t(246) = -4.20$, $p < .001$, but distinctive NoF did not, partial $r = -.02$, $t(246) = -0.27$, $p > .7$. After shared NoF entered the equation, the contribution of distinctive NoF again did not quite reach significance, partial $r = -.11$, $t(245) = -1.73$, $p > .08$. For lexical decision, shared NoF again significantly predicted decision latency, partial $r = -.25$, $t(246) = -4.04$, $p < .001$, but distinctive NoF did not, partial $r = -.11$, $t(246) = -1.80$, $p > .07$. After shared NoF entered the equation, however, both shared and distinctive NoF were significant predictors: shared NoF, partial $r = -.31$, $t(245) = -5.00$, $p < .001$; distinctive NoF, partial $r = -.21$, $t(245) = -3.40$, $p < .01$. These regression analyses confirm the results of the three studies. Both the number of shared features and, to a lesser extent, the number of distinctive features influence decision latencies.

Study 3 and the regression analyses also help to clarify the interpretation of Study 2. We stated above that the results of Study 2 were most likely due to a facilitative influence of shared features, but it was possible that they were due to an inhibitory influence of distinctive features. These regressions support the former conclusion, but not the latter. That is, they show that the greater the number of shared features, the shorter the decision latency in both tasks. Although the influence of distinctive NoF was small, that influence was facilitatory rather than inhibitory (i.e., the partial correlation with decision latency was negative rather than positive). Therefore, it can be concluded that the results of Study 2 were due to the strong positive effect of shared features.

One final aspect of Study 3 should be noted. To a large extent, semantic density of a concept depends on the number of shared features because feature correlations are calculated using shared features only. These measures can be decoupled in some circumstances (as they were in Study 2) because shared features do differ in the degree to which they are correlated with other features. However, because Study 3 included a condition, the low NoF shared-manipulated condition, that contained an extremely low number of shared features (3.1 on average), it was not possible to match it to the high NoF shared-manipulated condition on semantic density. Therefore, it might be possible that semantic density played a role in Study 3. That is, it could be the case that shared features that are highly correlated with other shared features are particularly important (i.e., not all shared features are created equal). This possibility does not change our basic conclusions regarding the importance of shared features, however, because semantic density was equated in Study 2 and an influence of shared features was obtained.

Study 4

In previous articles concerning NoF effects, there has been no discussion regarding the manner in which the types of featural knowledge may be differentially key to understanding these effects. Knowledge type (or feature type) analyses have been applied to other phenomena, notably category-specific semantic deficits. For example, Cree and McRae (2003) classified all of the features in McRae et al.'s (2005) norms in terms of what they called a brain-region taxonomy. They used research in neuroscience and cognitive neuropsychology to develop a taxonomy of types of knowledge that are at least somewhat separable in the brain. The taxonomy includes nine knowledge types: visual form and surface features (e.g., <has legs>, <made of metal>), tactile features (<is smooth>), functional features (<used by turning>), color (<is red>), visual motion features (<swims>), taste (<is sweet>), sound (<moos>), smell (<smells bad>), and encyclopedic features (<lives in Africa>).

In the present analyses, we investigated the knowledge types that exist in the concepts used in Studies 1–3 in terms of the number of shared and distinctive features within each class. Note that taxonomic features such as <is a vehicle> were excluded because they were excluded from all feature counts presented in this article.

In our studies, concreteness decision latency was a somewhat more sensitive measure than was lexical decision latency. Particularly in the case of concreteness decisions ("Is it a concrete object? That is, is it touchable?"), one would expect that features that strongly signal a concrete object, and perhaps in addition, are concrete in and of themselves, might be important. In fact, this might also apply to lexical decision, but in a somewhat muted manner. That is, if a letter string activates aspects of meaning that signal a concrete object, then that letter string must be a word.

Arguments could be made for the potential importance of basically all of the knowledge types. Table 7 presents the overall number of features of each type (i.e., shared plus

Table 7

Mean number per concept of shared plus distinctive features of each knowledge type

Knowledge type	Study 1		Study 2		Study 3			
	Low NoF	High NoF	Low shared	High shared	Shared manipulated		Distinctive manipulated	
					Low NoF	High NoF	Low NoF	High NoF
Visual form and surface	3.02	5.55	4.35	4.96	2.85	5.05	2.80	5.60
Encyclopedic	2.05	3.58	3.20	2.35	2.45	3.15	1.90	4.60
Tactile	0.44	0.58	0.42	0.56	0.15	0.55	0.90	0.40
Taste	0.11	0.16	0.11	0.13	0.05	0.05	0.05	0.05
Sound	0.08	0.34	0.20	0.20	0.15	0.35	0.00	0.05
Smell	0.03	0.09	0.07	0.00	0.10	0.00	0.15	0.20
Color	0.58	1.06	0.60	0.95	0.65	1.00	0.70	0.70
Visual motion	0.47	0.84	0.42	0.58	0.10	0.95	0.35	0.25
Function	2.25	3.52	3.29	2.64	3.00	4.05	2.55	3.55
Overall NoF	9.03	15.72	12.66	12.37	9.50	15.15	9.40	15.40

distinctive), whereas Table 8 presents shared features only, and Table 9 distinctive features only. Visual form and surface features should be central because they describe physical parts (<has a tail>), shapes (<is round>), or the materials used to make an object (<made of metal>), all of which unambiguously signal a concrete object. Furthermore, visual form and surface features are relatively plentiful in the norms. Paivio (1986) claimed that in addition to being able to verbally reason about concrete things, people can generate mental images for concrete words because they refer to physical things in the world that we perceive. He argued that this additional information associated with concrete words makes their mental representations richer and easier to activate. Presumably, the number of visual features, and particularly the number of shared (relatively common) visual features should increase ease of imageability and thus ease of a concreteness decision. That is, shared visual form and surface features represent parts and other aspects of living and nonliving things that are common to many things, and thus would be imaged with higher frequency relative to distinctive form and surface features.

Table 8

Mean number per concept of shared features of each knowledge type

Knowledge	Study 1		Study 2		Study 3			
	Low NoF	High NoF	Low shared	High shared	Shared manipulated		Distinctive manipulated	
					Low NoF	High NoF	Low NoF	High NoF
Visual form and surface	2.70	4.14	2.76	4.60	1.50	4.00	2.50	3.15
Encyclopedic	1.27	1.70	1.07	1.78	0.45	1.20	1.00	1.00
Tactile	0.41	0.48	0.35	0.55	0.05	0.30	0.85	0.25
Taste	0.09	0.11	0.11	0.11	0.00	0.10	0.05	0.00
Sound	0.08	0.14	0.11	0.15	0.05	0.00	0.00	0.00
Smell	0.00	0.05	0.02	0.00	0.00	0.00	0.10	0.05
Color	0.58	1.02	0.56	0.93	0.55	1.00	0.70	0.70
Visual motion	0.36	0.52	0.25	0.51	0.00	0.70	0.20	0.15
Function	1.44	1.64	1.42	2.02	0.45	1.85	1.10	0.95
Shared NoF	6.93	9.80	6.65	10.65	3.05	9.15	6.50	6.25

Table 9

Mean number per concept of distinctive features of each knowledge type

Knowledge type	Study 1		Study 2		Study 3			
	Low NoF	High NoF	Low shared	High shared	Shared manipulated		Distinctive manipulated	
					Low NoF	High NoF	Low NoF	High NoF
Visual form and surface	0.31	1.41	1.58	0.36	1.30	1.05	0.30	2.45
Encyclopedic	0.78	1.88	2.13	0.58	2.00	1.95	0.90	3.60
Tactile	0.03	0.09	0.07	0.02	0.10	0.25	0.05	0.15
Taste	0.02	0.06	0.02	0.02	0.05	0.00	0.00	0.05
Sound	0.00	0.20	0.09	0.05	0.10	0.35	0.00	0.05
Smell	0.03	0.05	0.05	0.00	0.10	0.00	0.05	0.15
Color	0.00	0.05	0.04	0.02	0.10	0.00	0.00	0.00
Visual motion	0.11	0.33	0.16	0.07	0.15	0.25	0.15	0.10
Function	0.81	1.88	1.89	0.60	2.55	2.20	1.45	2.60
Distinctive NoF	2.09	5.95	6.03	1.72	6.45	6.05	2.90	9.15

Tactile features may also be central because if someone knows what an object feels like, it must be touchable. Thus, tactile features should facilitate concreteness decisions in particular. Functional knowledge also seems like it should predict decision latencies because if someone has physically used an object for some purpose (like a fork for eating), then it must be a touchable object.

Other sensory features might also be important. For example, color likely signals something that is touchable, although there are salient exceptions such as the sky. Taste and smell features are somewhat ambiguous. Concrete objects such as foods and fruits and vegetables have salient taste features, and sometimes smell features as well. However, gases seem like they have a taste, and definitely have a smell, but they are not touchable. Sounds generally emanate from concrete objects, although a sound such as thunder is a salient exception, and sound itself is not touchable or concrete. Visual motion features might signal a concrete object because if a person has seen something move about on its own, it is probably touchable. Some naturally occurring non-concrete things such as clouds would be salient exceptions.

Finally, the degree to which encyclopedic features are related to concreteness is less clear. When Cree and McRae (2003) originally classified features using the brain-region knowledge type taxonomy, features that were not clearly part of the other knowledge types were classified as encyclopedic. Thus, encyclopedic features include somewhat of a mixed bag of information types. Some of these features, such as <has sentimental value> or <is fun>, are clearly not related to concreteness. However, many others do signal a concrete object. For example, some encyclopedic features describe characteristic behaviors of entities, such as <lays eggs> or <hibernates>. A large number of encyclopedic features convey information about location and time, such as <lives in water>, <grows in gardens>, <used on farms>, and <worn in winter>. A conservative estimate is that these types of features comprise at least 60% of the encyclopedic features when shared and distinctive features are combined. These features signal two types of information. First, the activities that are part of these featural descriptions carry information about concreteness in that

they signal that the thing is alive, grows, is used, or is worn. Second, they carry contextual information regarding the situations in which the object or entity tends to occur, such as the location at which it tends to be found. Bransford and McCarrell (1974) have made claims regarding contextual information that are relevant to these encyclopedic features. They argued that because people interact directly with concrete things in the world, but not with abstract concepts, concrete words have more contextual information associated with them (such as place and time as indexed by these encyclopedic features), facilitating the computation and use of these concepts. It is possible that the greater the amount of this contextual information associated with a concept, the easier a concreteness decision may be (and perhaps a lexical decision as well). In addition, common contexts that are shared by numerous objects and entities may play a particularly important role.

Methods

We counted the number of shared and distinctive features of each knowledge type in each concept used in Study 1, 2, or 3. These counts were the independent variables of interest in stepwise regression analyses. The dependent variables were concreteness and lexical decision latencies when the items from all three studies were combined. As in the regression analyses presented above, there were 250 data points for each task, and $\ln(\text{BNC})$ word frequency and length in letters were forced in on the first step to account for basic word-reading non-semantic variables.

Results and discussion

The partial correlations and significance values are presented in Table 10. Note that tolerances were sufficiently high in all cases. For concreteness decision latency, the knowledge types that initially significantly predicted decision latency were shared encyclopedic features and shared visual form and surface features. Due to partitioning of variance, however, four variables entered the regression equation in the following order: shared encyclopedic features,

Table 10
Predicting decision latency with knowledge types

Predictor	Partial <i>r</i>	<i>t</i> (243)	<i>p</i>
<i>Concreteness decision</i>			
Shared encyclopedic	-.23	-3.61	.001
Shared tactile	-.16	-2.51	.01
Shared visual form and surface	-.18	-2.88	.004
Shared taste	-.14	-2.16	.03
<i>Lexical decision</i>			
		<i>t</i> (243)	<i>p</i>
Shared encyclopedic	-.20	-3.12	.002
Shared visual form and surface	-.19	-2.93	.004
Shared tactile	-.16	-2.50	.01
Distinctive visual form and surface	-.16	-2.51	.01
Shared taste	-.14	-2.12	.04

shared tactile features, shared visual form and surface features, and shared taste features.

The results were similar for lexical decision latency. The knowledge types that initially predicted decision latency were again shared encyclopedic and shared visual form and surface features. Due to partitioning of variance, five variables entered the equation in the following order: shared encyclopedic, shared visual form and surface, shared tactile, distinctive visual form and surface, and shared taste features.

These regressions further support the conclusion that shared features are central to NoF effects. Significant predictors consisted almost exclusively of shared features. Shared visual form and surface features are important because parts, shape, and materials are highly informative for object concepts, and because features corresponding to this knowledge type are relatively numerous. The one case in which a distinctive knowledge type predicted variance in decision latencies involved visual form and surface knowledge. Shared encyclopedic features predicted decision latencies in both tasks. As stated above, although these features do not fit into the other brain-region knowledge types, many are central to object concepts and often do signal concreteness (such as where something lives or grows, its characteristic behaviors, or where an object tends to be located or used). In addition, many of the encyclopedic features in the norms correspond to the amount of contextual (situational) knowledge that people possess about concrete objects. Approximately 70% of the shared encyclopedic features were of this sort, versus about 60% overall for encyclopedic features. As with visual form and surface features, these types of encyclopedic features are produced frequently by participants in the norming task.

With respect to shared tactile features, they predicted concreteness decision latency presumably because the instructions asked participants to indicate whether each word referred to something that is a touchable. However, it was somewhat surprising that they predicted lexical decision latency as well, particularly because tactile features are not numerous. It is perhaps the case that it is meaningful if any tactile features are produced for a concept. That is, if a person has actually handled something and thus knows what it feels like, they are likely to produce only one tactile feature such as <feels rough>, and this

information signals that a letter string corresponds to a word that refers to a concrete touchable object. Therefore, although tactile features are not numerous, their presence is meaningful, signaling personal contact with the object or entity. Finally, the amount of information about how objects taste played a role. Taste features are important for fruits, vegetables, and foods in the present experiments. In other words, if people have eaten something, then they know it is concrete, and they know that its name is a word. As is the case with tactile features, people are likely to produce only one taste feature per concept (e.g., <tastes sweet>).

Finally, it is somewhat surprising that functional knowledge did not predict decision latencies, particularly given the centrality of function in research on concepts and semantic memory, and the relatively large number of functional features in the norms. Knowledge regarding how something is used seems like it should be a reliable cue to concreteness or whether something is touchable in particular. It is not entirely clear why the number of functional features failed to predict either concreteness or lexical decision latency. One possibility is that the null effects of the number of function features hinges on the fact that the number of different functions that an object performs is not informative of concreteness, at least when there exists at least one known function. That is, according to the norms, some objects are used for multiple related functions, such as <used by bouncing>, <used by throwing>, <used for sports>, and <used for games> for *ball*, whereas others have one feature referring to a single dominant function, as in <used for cooking> for *pot*. It may be the case that the fact that an object is consistently used for a single function is as salient and informative a cue as is the fact that it is used for multiple related functions. Also note that for a feature to be part of the counts that were used in this (and our other related) research, a minimum of five of thirty participants must have produced it. Sometimes, as was the case with *pot* <used for cooking>, one functional feature dominated (it was produced by 28 of 30 participants). Therefore, other related functions did not reach threshold. For example, with *pot*, three participants produced <used for boiling water> and two produced <used for making stew>. As discussed by Barsalou, Sloman, and Chaigneau (2005), some features like <used for cooking> may stand in for a number of related functions. It may be the case that functional feature counts on a concept-wise basis might be somewhat obscured in some of these cases, which may have contributed to the null effect of the number of functional features.

In summary, the types of information that most strongly drive NoF effects are knowledge regarding objects' common parts, shapes, and materials, where and when they commonly live, grow, and are located, how they feel, and how they taste. These types of knowledge influence both deciding that a word refers to a concrete object, and deciding that a concept's name is indeed a word. Finally, these analyses buttress our conclusions regarding the importance of shared features because virtually all of the significant predictors involve shared features of various types.

General discussion

Pexman et al. (2002, 2003) found that concepts with many features were responded to more quickly than those with few features, and this was taken as evidence of a computational advantage for words rich in semantic representation. The present Study 1 was essentially a replication of Pexman's studies, showing that the NoF effect is robust when using larger sets of items that were better equated on a large number of variables. The remainder of the present article adds to the understanding of these empirical phenomena in two major ways. First, we built on previous research that has demonstrated that the distinction between shared and distinctive features has implications for a number of aspects of conceptual processing. Studies 2 and 3 showed that increasing the number of shared features facilitates processing to a greater extent than does increasing the number of distinctive features, although both types of features do facilitate decision latencies. Second, previous research has provided a great deal of evidence that object concepts are represented over multiple modality-specific brain regions. We used insights from these studies, most particularly the brain-region feature type taxonomy of Cree and McRae (2003), to test whether types of knowledge that are associated with various brain regions contribute differentially to NoF effects. Study 4 showed that NoF effects are due primarily to shared encyclopedic, visual form and surface, tactile, and taste features. In the next section, we describe how our results and analyses constrain interpretations of NoF effects.

Decisions versus computing a concept

Much of the previous discussion has been framed in terms of a decision-making account of our results. That is, the idea has been that certain types of knowledge signal concrete objects, and features that are common to numerous concrete objects are better cues to concreteness and lexical decisions than are those that apply to very few objects.

One might wonder if it is possible to account for the present results, at least those in the concreteness decision task, by focusing instead on the computation of word meaning. We begin with the assumption that a model based on distributed semantic representations is required. Consider, for example, distributed attractor network models of semantic memory such as those used by Cree, McRae, and McNorgan (1999) or Plaut and Shallice (1993). Typically, decision latency is simulated in these models using some metric of network error, such as mean squared error or cross-entropy. Error is recorded as the network settles to a stable representation. Distributed feature-based models of semantic memory have sparse semantic representations because each concept includes only a small subset of the total number of possible features. Therefore, the model has a bias to turn semantic units off because it is trained to turn each unit off for the majority of concepts. Because the most challenging aspect of the model's computations corresponds to activating over time those feature units that are part of a concept, by far the bulk of the error as the network settles is asso-

ciated with these units. To correctly simulate NoF effects, a model would have to show less error for concepts that contain a greater number of semantic feature units that need to be activated. Therefore, in a system such as this, all other factors being equal (as they were in our studies), concepts with many active units (high NoF) engender greater error than those with few active units (low NoF) as the network settles. Therefore, it seems unlikely that any model of this sort that uses an error measure to simulate latencies would be able to produce a processing advantage for high NoF concepts.

To examine this issue, we implemented a model in which the words representing each of the 541 concepts in the norms were represented as three randomly activated units out of 30 word form input units. There were 2349 semantic feature output units, each representing a feature from McRae et al.'s (2005) norms (all taxonomic features were excluded, as in all analyses reported above). The model was trained to settle on a concept's representation over 20 iterations ("time steps") using the continuous recurrent backpropagation-through-time algorithm (Pearlmutter, 1995) and cross-entropy as the error metric. The identical model is described in full in Cree et al. (2006). Following training, we simulated Study 1 using the items from that study. As expected, error over the semantic units was greater for high NoF concepts than for low NoF concepts until the last few time steps when both settled to approximately the same error levels. There was no difference in cross-entropy error between the high and low-shared conditions of Study 2 (these conditions were identical in terms of overall NoF). Finally, as in the simulation of Study 1, the model incorrectly predicted reverse NoF effects in Study 3.

It is possible, of course, that other measures from a distributed attractor network could simulate decision latencies. One such measure might be the total amount of activation in the system, with the assumption being that decision latency is monotonically related to total activation. Again, the influence of such a measure is self-evident for basic NoF effects, and in this case, it correctly accounts for the results. In the simulation of Study 1, the total activation was greater for the high NoF group from time steps 10–20 of the computation of the semantic representations from word form. That is, the greater the number of features that are activated, the greater the overall activation, and the shorter the predicted decision latency.

The next question is whether a model of this sort could account for the influence of shared features. The answer appears to be no. The total activation measure shows extremely small differences between the high and low-shared groups of Study 2. For Study 3, the simulation predicts shorter decision latencies for the two high NoF conditions than for the two low NoF conditions. It does predict a small disadvantage late in processing (time steps 14–20) for the low NoF shared-manipulated condition as compared to the low NoF distinctive-manipulated condition, although this difference is much smaller than between the low NoF distinctive-manipulated condition and the two high NoF conditions, which does not mirror the results of Study 3. Thus, it does not appear that an attrac-

tor network model with distributed feature-based representations successfully accounts for the present results.

It should be noted that a failure to simulate human empirical results is actually a null effect, similar to an experimental null effect. That is, it is possible that a different type of distributed network might be able to simulate our studies. However, given these simulation results, one question that follows is whether it is possible that a decision-making account could provide an explanation for these results. There exist other studies of distinctive versus shared features that highlight the importance of decision processes. Contrast the present results with those of Cree et al. (2006). They found advantages in feature verification latencies for distinctive over shared features in both living and nonliving things (although see Randall et al., 2004, who found an advantage for distinctive features of nonliving things, but a disadvantage for distinctive features of living things). At first blush, these results appear to be inconsistent with those of the present Studies 2 and 3. In Study 1 of Cree et al., a concept name was presented first, followed by a feature name. Participants were faster to decide that a feature was part of a concept if the feature was distinctive. That is, there was an advantage for distinctive features when the decision entailed evaluating a specific concept-feature relation, a type of decision that differs substantially in nature from deciding whether a letter string corresponds to a concrete concept or an English word.

Cree et al. (2006) simulated this result by inputting the concept name, then recording the activation of the target feature node (using the same model as used herein). Distinctive features were activated more quickly than were matched shared features. Given that the decision entailed directly evaluating a concept-feature relation, using the assumption that the network's activation of a specific feature is monotonically related to feature verification latency provided a good fit to the data. Thus, the difference in the importance of distinctive versus shared features in the two studies is due directly to the decision that was required, and this difference is apparent in the ability of an attractor network to simulate the data.

One possible method to simulate a concrete-abstract decision process is to use something akin to a random walk model of conceptual decision making (Joordens, Piercey, & Azerbehi, 2003; Ratcliff, Gomez, & McKoon, 2004). Such a model includes two thresholds, one representing a concrete decision and another representing an abstract decision. The system might behave as follows. A word is presented and its semantic features are computed. When a feature that signals a concrete object is activated, such as a visual form and surface feature, certain encyclopedic features, tactile features, or taste features, the system moves closer to the concrete threshold. The greater the number of these types of features that are computed, the shorter the decision latency. Shared features that are common to numerous concrete objects would thus facilitate decisions for two reasons. First, they are better cues to concreteness because they better cue the fact that the word refers to something that is a member of the large category of concrete objects and entities. Second, they are more common (frequent) familiar in of themselves. Therefore, they are cues that people are more familiar with.

Similar well-known results have been found for category verification tasks, such as "Is a robin a bird?", that are analogous to the concreteness decision task used in our studies. In category verification, the relevant shared features are those that are shared by members of the category, that is, *bird* in this case (Rosch & Mervis, 1975). Typical exemplars such as *robin* for the category *bird* possess a greater number of features that are shared by category exemplars than do atypical exemplars such as *ostrich*. As such, participants rate those exemplars as more typical, and they verify them more quickly as a member of the category (Smith et al., 1974). These category verification effects have most often been explained in terms of decision processes rather than the speed of concept retrieval. Note that the notion of shared features is thus relative to the required decision. That is, what is shared across concepts depends critically on the category that is relevant to the task.

If our concreteness decision results are analogous to category verification, one would expect that concreteness ratings might mirror them in the same manner that typicality ratings mirror category verification latency. That is, if off-line concreteness ratings show a pattern that is similar to the concreteness decision latencies, then this would provide additional evidence for our decision-making account. We tested this idea by collecting concreteness ratings for the 250 stimuli included in the three studies from 31 new participants using instructions based on those of Toggia and Battig (1978). As is customary for concreteness ratings, a 7-point scale was used with seven corresponding to definitely concrete, and the resulting means were scaled by 100. An additional 250 words referring to abstract concepts were included to anchor the scale. As would be expected given that all of our items refer to concrete objects and entities, the variance in the ratings for those items was quite low and ratings tended to be in the upper end of the scale. Although the differences were quite small, the results did mirror the concreteness decision latencies. For Study 1, concreteness ratings were higher for high NoF ($M = 634$) than for low NoF items ($M = 614$), $t(126) = 3.07, p < .01$. For Study 2, concreteness ratings were higher for high shared ($M = 632$) than for low-shared items ($M = 621$), $t(108) = 2.20, p < .03$. For Study 3, the pattern of means was similar to the concreteness decision latencies in that concreteness ratings were lowest for the low NoF shared-manipulated condition ($M = 610$), and the other three conditions were similar (high NoF shared manipulated, $M = 629$; low NoF distinctive manipulated, $M = 623$; high NoF distinctive manipulated, $M = 622$). However, neither of main effects nor the interaction were significant, $p > .1$ in all three cases. In summary, the concreteness ratings pattern with the concreteness decision latencies, thus bolstering our decision-making account.

In category verification tasks, features that are not part of category exemplars cue a "no" response. Research has begun to be conducted on abstract concepts, and it is clear that the types of knowledge underlying them differ from that of concrete concepts (Barsalou & Wiemer-Hastings, 2005). Therefore, the types of knowledge computed when a word referring to an abstract concept is read differ from concrete concepts, and they push the system to-

ward making an abstract “no” decision. When an abstract or intangible feature is activated, the system moves closer to the abstract threshold. Whatever criterion is reached first determines the system’s decision. The fact that abstract decision latencies were longer than concrete decision latencies in all three experiments may be due to the fact that knowledge underlying abstract concepts is less clear. More likely, this occurred because the concreteness decision task was presented to participants as if it required a yes/no response, and negative responses are virtually always slower than positive ones in any binary decision task. In fact, one way that such “no” responses could be modeled is to include an inherent drift toward that decision (in this case, that the word refers to an abstract concept). That is, there could be an inherent drift toward a “no” response so that positive evidence is required to respond “yes”, and a “no” response is generated in the absence of such evidence. This is similar in spirit to the time deadline used in Grainger and Jacobs’ (1996) multiple read-out model.

Another way to think about the influence of shared features on concreteness decision is to consider analogous orthographic effects on lexical decision. Numerous studies have shown that orthographic factors such as frequent letters or sub-lexical letter strings facilitate lexical decisions (Coltheart et al., 1977). For example, a variable such as orthographic neighborhood size indexes the degree to which a letter string shares orthographic “features” (letters, bigrams, trigrams, etc.) with other words. Just as being more similar to other words facilitates lexical decisions, being more similar to other concrete objects and entities facilitates concreteness decisions.

With respect to our lexical decision results, we controlled for word recognition variables that are known to influence lexical decision latencies. Therefore, any differences found in our studies were due to semantic variables because the words in various conditions were equally word-like with respect to orthographic and phonological variables. Semantic knowledge does influence lexical decisions, although not as strongly as it influences semantic-based decisions (Becker et al., 1997). When making a decision regarding whether a letter string is a word, information signaling that it refers to a concrete object also signals that the letter string is indeed a word. Of course, the same could be said for numerous types of information; semantic knowledge that signals an abstract concept, an activity (a verb such as *jog*), and so on, also signals that the letter string corresponds to a word. In our studies, however, we were concerned with differences within concrete concepts only. Therefore, features that are shared by numerous concrete concepts push a binary decision system toward a “word” decision, but presumably their influence would be weaker (less easily detectable) than in the case in which the decision was indeed whether the word referred to a concrete concept. Thus, the observed influence of shared features was weaker in lexical decision than in the concreteness decision tasks of Studies 2 and 3.

In summary, it is clear that the influence of a variable such as the degree to which features are shared across con-

cepts depends crucially on the task under consideration. Feature verification shows an advantage for distinctive features. Shared features inhibit a task such as picture naming in which participants must distinguish a basic level concept from other similar concepts. Finally, features that are shared across category members facilitate decisions regarding membership in that category.

Conclusion

The present research provides additional insight into why words that possess rich semantic representations are responded to more quickly. Using tightly matched lists of concrete concepts, we demonstrated that concepts with many features are responded to faster than those with few features in lexical and concreteness decision tasks. We also demonstrated that this facilitation is modulated by whether features are distinctive or shared. Finally, the most important types of knowledge driving these decision-based effects are shared visual form and surface, encyclopedic, tactile, and taste features.

Appendix A

Stimuli used in Studies 1A and 1B

Low number of features	High number of features
ball	axe
bayonet	balloon
bedroom	barn
birch	basement
biscuit	bathtub
bluejay	bear
boat	blender
broccoli	bra
brush	bus
bucket	camel
cabinet	camisole
catfish	cannon
cherry	canoe
clamp	carrot
cloak	chimp
cod	cougar
cupboard	cow
dish	coyote
doll	crab
dove	desk
eel	dog
finch	drapes
flamingo	fawn
garage	freezer
hook	garlic
hornet	goat
inn	gorilla
mackerel	grenade
mirror	hammer
mixer	hose
moth	iguana
oak	kettle

Appendix A (continued)

Low number of features	High number of features
otter	lettuce
pan	lion
parka	missile
peas	mushroom
pheasant	necklace
pin	nylons
pine	olive
platypus	pearl
pot	pen
pumpkin	pencil
razor	pickle
rhubarb	pig
rock	pony
rocker	potato
scissors	radish
shawl	rake
skillet	rat
skunk	robe
snail	rooster
spinach	sandals
squid	saucer
starling	screws
stone	seal
taxi	spear
tent	spoon
toilet	swimsuit
toy	sword
tray	tiger
trolley	toad
turnip	train
veil	tricycle
whale	wasp

Appendix B

Stimuli used in Studies 2A and 2B

Low number of shared features	High number of shared features
basement	bag
beaver	ball
bike	bazooka
bolts	bear
bomb	beetle
broom	bookcase
bull	bowl
calf	buffalo
chair	buggy
church	bureau
cigar	cabinet
cork	car
crown	caribou
curtains	coyote
dog	crab

Appendix B (continued)

Low number of shared features	High number of shared features
doll	dish
drapes	dress
faucet	dresser
flea	elk
garage	flute
gate	fox
gloves	frog
grenade	jacket
hornet	knife
jar	lion
jeans	lobster
kettle	marble
lime	mushroom
medal	oak
minnow	olive
mittens	pear
necklace	piano
octopus	plate
onions	pliers
pajamas	pony
pie	pot
pigeon	raccoon
pillow	radish
pine	scooter
prune	shack
pumpkin	shrimp
ruler	skillet
ship	slippers
sink	spatula
skirt	spear
skis	squirrel
sofa	stone
swimsuit	tack
tomahawk	tractor
train	tray
unicycle	tricycle
veil	truck
vine	tuba
violin	van
walrus	walnut

Appendix C

Stimuli used in Studies 3A and 3B

Low number of distinctive features	High number of distinctive features
<i>Distinctive features manipulated</i>	
brick	anchor
cedar	apartment
flamingo	barn
garage	basement
goldfish	bathtub

(continued on next page)

Appendix C (continued)

Low number of distinctive features	High number of distinctive features
hook	book
mirror	bra
parsley	canoe
rock	cigarette
sack	crowd
saddle	faucet
scissors	flea
seaweed	hose
shawl	kettle
shelves	nylons
skunk	onions
slippers	pyramid
stick	ruler
toy	sandals
worm	shovel
<i>Shared features manipulated</i>	
bagpipe	apron
beehive	balloon
bouquet	banner
cape	bolts
certificate	box
clock	carpet
envelope	cottage
fence	dolphin
hut	drapes
key	freezer
menu	goat
muzzle	pigeon
pepper	potato
pier	prune
pine	revolver
projector	screws
radio	sheep
tape	skirt
unicycle	sled
wheel	stool

References

Barsalou, L. W., Sloman, S. A., & Chaigneau, S. E. (2005). The HIPE theory of function. In L. Carlson & E. van der Zee (Eds.), *Representing functional features for language and space: Insights from perception, categorization and development* (pp. 131–147). New York: Oxford University Press.

Barsalou, L. W., & Wiemer-Hastings, K. (2005). Situating abstract concepts. In D. Pecher & R. Zwaan (Eds.), *Grounding cognition: The role of perception and action in memory, language, and thought* (pp. 129–163). New York: Cambridge University Press.

Becker, S., Moscovitch, M., Behrmann, M., & Joordens, S. (1997). Long-term semantic priming: A computational account and empirical evidence. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 23, 1059–1082.

Binder, J. R., Westbury, C. F., McKiernan, K. A., Possing, E. T., & Medler, D. A. (2005). Distinct brain systems for processing concrete and abstract concepts. *Journal of Cognitive Neuroscience*, 17, 1–13.

Borowsky, R., & Mason, M. E. J. (1996). Semantic ambiguity effects in word identification. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22, 63–85.

Bransford, J. D., & McCarrell, N. S. (1974). A sketch of a cognitive approach to comprehension: Some thoughts on what it means to comprehend. In W. Weimer & D. Palermo (Eds.), *Cognition and the symbolic processes* (pp. 189–230). Hillsdale, NJ: Erlbaum.

Bueno, S., & Frenck-Mestre, C. (2008). The activation of semantic memory: Effects of prime exposure, prime-target relationship, and task demands. *Memory & Cognition*, 36, 882–898.

Burnard, L. (2000). *British National Corpus user reference guide Version 2.0*. Retrieved May 20, 2002, from British National Corpus, Oxford University Computing Service. Available from <http://hcu.ox.ac.uk/BNC/World/HTML/urg.html>.

Cohen, J. D., MacWhinney, B., Flatt, M., & Provost, J. (1993). PsyScope: A new graphic interactive environment for designing psychological experiments. *Behavior Research Methods*, 25, 257–271.

Collins, A. M., & Quillian, M. R. (1969). Retrieval time from semantic memory. *Journal of Verbal Learning and Verbal Behavior*, 8, 240–247.

Coltheart, M., Davelaar, E., Jonasson, J. T., & Besner, D. (1977). Access to the internal lexicon. In S. Dornic (Ed.), *Attention and performance* (Vol. 6, pp. 535–555). Hillsdale, NJ: Erlbaum.

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.

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

Davis, C. J. (2005). N-Watch: A program for deriving neighborhood size and other psycholinguistic statistics. *Behavior Research Methods*, 37, 65–70.

de Groot, A. M. (1989). Representational aspects of word imageability and word frequency as assessed through word association. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 15, 824–845.

Garrard, P., Lambon Ralph, M. A., Hodges, J. R., & Patterson, K. (2001). Prototypicality, distinctiveness, and intercorrelation: Analyses of the semantic attributes of living and nonliving concepts. *Cognitive Neuropsychology*, 18, 125–174.

Goldberg, R. F., Perfetti, C. A., & Schneider, W. (2006). Perceptual knowledge retrieval activates sensory brain regions. *Journal of Neuroscience*, 26, 4917–4921.

Grainger, J., & Jacobs, A. M. (1996). Orthographic processing in visual word recognition: A multiple read-out model. *Psychological Review*, 103, 518–565.

Hino, Y., & Lupker, S. J. (1996). Effects of polysemy in lexical decision and naming: An alternative to lexical access accounts. *Journal of Experimental Psychology: Human Perception and Performance*, 22, 1331–1356.

Hino, Y., Pexman, P. M., & Lupker, S. J. (2006). Ambiguity and relatedness effects in semantic tasks: Are they due to semantic coding? *Journal of Memory and Language*, 55, 247–273.

Humphreys, G. W., Riddoch, M. J., & Quinlan, P. T. (1988). Cascade processes in picture identification. *Cognitive Neuropsychology*, 5, 67–103.

James, C. T. (1975). The role of semantic information in lexical decisions. *Journal of Experimental Psychology: Human Perception and Performance*, 1, 130–136.

Jastrzemski, J. E. (1981). Multiple meanings, number of related meanings, frequency occurrence, and the lexicon. *Cognitive Psychology*, 13, 278–305.

Joordens, S., Piercey, C. D., & Azerbehi, R. (2003). From word recognition to lexical decision: A random walk along the road of harmony. In: F. Detje, D. Dörner, & H. Schaub (Eds.), *Proceedings of the fifth international conference on cognitive modeling* (pp. 141–146).

Kellas, G., Ferraro, F. R., & Simpson, G. B. (1988). Lexical ambiguity and the time course of attentional allocation in word recognition. *Journal of Experimental Psychology: Human Perception and Performance*, 14, 601–609.

Kučera, H., & Francis, W. N. (1967). *Computational analysis of present-day American English*. Providence, RI: Brown University Press.

Laws, K. R., & Gale, T. M. (2002). Category-specific naming and the visual characteristics of line drawn stimuli. *Cortex*, 38, 7–21.

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

- Masson, M. E. J. (1995). A distributed memory model of semantic-priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21, 3–23.
- Masson, M. E. J., & Loftus, G. R. (2003). Using confidence intervals for graphically based data interpretation. *Canadian Journal of Experimental Psychology*, 57, 203–220.
- McNamara, T. P. (1992). Priming and constraints it places on theories of memory and retrieval. *Psychological Review*, 99, 650–662.
- 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*, 37, 547–559.
- 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.
- Medin, D. L. (1989). Concepts and conceptual structure. *American Psychologist*, 44, 1269–1481.
- Myers, J. L. (1979). *Fundamentals of experimental design*. Boston, MA: Allyn and Bacon.
- Page, M. (2000). Connectionist modeling in psychology: A localist manifesto. *Behavioral and Brain Sciences*, 23, 443–512.
- Paivio, A. (1986). *Mental representations: A dual coding approach*. New York: Oxford University Press.
- Pearlmutter, B. A. (1995). Gradient calculation for dynamic recurrent neural networks: A survey. *IEEE Transactions on Neural Networks*, 6, 1212–1228.
- Pecher, D., Zeelenberg, R., & Barsalou, L. W. (2003). Verifying different-modality properties for concepts produces switching costs. *Psychological Science*, 14, 119–124.
- Pexman, P. M., & Hope, J. (in preparation). Is more always better? Dissecting number of features effects and semantic processing.
- Pexman, P. M., Holyk, G. G., & MonFils, M.-H. (2003). Number-of-features effects and semantic processing. *Memory & 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.
- Plaut, D. C., & Shallice, T. (1993). Deep dyslexia: A case study of connectionist neuropsychology. *Cognitive Neuropsychology*, 10, 377–500.
- Randall, B., Moss, H. E., Rodd, J., Greer, M., & Tyler, L. K. (2004). Distinctiveness and correlation in conceptual structure: Behavioral and computational studies. *Journal of Experimental Psychology: Language, Memory, and Cognition*, 30, 393–406.
- Ratcliff, R., Gomez, P., & McKoon, G. (2004). A diffusion model account of the lexical-decision task. *Psychological Review*, 111, 159–182.
- Riddoch, M. J., & Humphreys, G. W. (1987). Picture naming. In G. W. Humphreys & M. J. Riddoch (Eds.), *Visual object processing: A cognitive neuropsychological approach* (pp. 107–143). London: Erlbaum.
- Rips, L. J., Shoben, E. J., & Smith, E. E. (1973). Semantic distance and the verification of semantic relations. *Journal of Verbal Learning and Verbal Behavior*, 12, 1–21.
- Rosch, E., & Mervis, C. B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology*, 7, 573–605.
- Rubenstein, H., Garfield, L., & Millikan, J. A. (1970). Homographic entries in the internal lexicon: Effects of systematicity and relative frequency of meanings. *Journal of Verbal Learning and Verbal Behavior*, 10, 57–62.
- Simmons, K. W., & Barsalou, L. (2003). The similarity-in-topography principle: Reconciling theories of conceptual deficits. *Cognitive Neuropsychology*, 20, 451–486.
- Smith, E. E., Shoben, E. J., & Rips, L. J. (1974). Structure and process in semantic memory: A feature model for semantic decisions. *Psychological Review*, 81, 214–241.
- Strain, E., Patterson, K., & Seidenberg, M. S. (1995). Semantic effects in single-word naming. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21, 1140–1154.
- Toglia, M. P., & Battig, W. R. (1978). *Handbook of semantic word norms*. New York: Erlbaum.