

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Which semantic properties of a feature affect access to an object concept?

Permalink

<https://escholarship.org/uc/item/841029b4>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 44(44)

Authors

Antal, Caitlyn
de Almeida, Roberto G.
Johns, Brendan

Publication Date

2022

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

Which semantic properties of a feature affect access to an object concept?

Caitlyn Antal (caitlyn.antal@mail.mcgill.ca)
Department of Psychology, McGill University
Montreal, QC, H3A 1G1

Roberto G. de Almeida (roberto.dealmeida@concordia.ca)
Department of Psychology, Concordia University
Montreal, QC, H4B 1R6

Brendan T. Johns (brendan.johns@mcgill.ca)
Department of Psychology, McGill University
Montreal, QC, H3A 1G1

Abstract

We investigated how the properties of lexical items, which label object features, affect concept tokening. We addressed this issue by modeling data from three sources: (1) norms obtained from a dataset of 78,000 features to a set of pictures representing living and nonliving objects; (2) accuracy data from a picture-word priming congruency task with stimuli presented for 50-60 milliseconds; and (3) corpus data on the lexical properties of four different social usage count measures. We conducted two sets of analyses: one relying on sample count-based measures (i.e., measures based on the norming study: sample frequency, cue validity, feature distinctiveness), and a second relying on the social usage count-based measures (i.e., word frequency (WF), contextual diversity (CD), discourse contextual diversity (DCD), and user contextual diversity (UCD)). Contrasting count and social usage-based measures allowed us to gain insight into the contribution of diverse semantic and socially oriented contextual measures of lexical items, and how they may affect concept tokening. Our results show that cue validity and feature distinctiveness were negative predictors of participants' accuracy to congruency decisions—an effect which was more pronounced for distinctive features of living things. There was also a noticeable advantage for the UCD and DCD variables, over CD and WF. Overall, our results suggest that the conceptual system may be organized as a function of both, intrinsic properties of object features and usage based contextual measures of lexical items that label these features.

Keywords: concepts; objects; features; prototype theory; word-picture congruency; lexical frequency; contextual diversity.

Introduction

What is in a concept? And how are concepts attained through vision and language? It is now virtually a consensus that the content of a concept is, to a large extent, determined by the features of the concept's referent (pace atomism; see Fodor & Pylyshyn, 2015). To wit, the idea is that the content of a concept such as DOG is determined by its purported constitutive features—the likes of *bark*, *fur*, *pet*, and *canine*—whether they are *necessary* or not (e.g., Smith & Medin, 1981; Moss, Tyler, & Taylor, 2007; Tyler & Moss, 2001; Rogers & McClelland, 2004). But what role does a feature play in accessing an object concept? If the feature

bark is indeed a constituent of the concept DOG, does recognition of the word “bark” provide access to the content of DOG? And to what extent do properties of a lexical label—viz., the word “bark” for the feature *bark*—influence access to its host concept DOG?

In order to address these questions, we modeled data from three main sources: (1) norms obtained from a dataset of 78,000 features produced by 100 participants (Antal & de Almeida, 2022a) who were presented with a set of pictures representing living and nonliving objects (Snodgrass & Vanderwart, 1980); (2) accuracy data from 71 participants who performed a picture-word priming congruency (PWPC) task with stimuli presented for 50-60 milliseconds (Antal & de Almeida, 2022b) and (3) corpus data on the lexical properties of two types of features: salient (i.e., those that have greater probability of occurrence in concepts, based on the norming task) and non-salient ones (i.e., features selected based on their low-probability of occurrence). The corpora data were obtained relying on four methods: (a) word frequency (WF), (b) contextual diversity (CD; see Adelman, Brown, & Quesada, 2006), (c) user contextual diversity (UCD; Johns, 2021a), and (d) discourse contextual diversity (DCD; Johns, 2021a). Data from WF, CD, UCD and DCD were obtained from 334,345 Reddit users who produced at least 3,000 comments each, across 30,327 subReddits, yielding approximately 55.7 billion words (Johns, 2021a).

Our investigation was motivated by the purported key role that features of concepts may play in concept tokening (that is, in the process of accessing the *content* of a concept). Consider, for instance the process of concept tokening depicted in Figure 1. The figure represents the two main ways in which a concept can be tokened perceptually—i.e., visually, upon seeing a referent, and linguistically, via its lexical label. The products of both, object recognition and word recognition systems, in principle, make contact with similar content, at a minimum the concept DOG. But if *bark* is a salient feature of the concept DOG, does it provide access to DOG after it's lemma (in principle “bark”) is processed? We reasoned that if the content of a concept is in large part determined by its constituent features, a lexical label representing its most salient feature should also yield access to its associated concept. This prediction is embedded in a

host of concept theories that rely on feature weights, as well as the strengths of links in a network between concepts or between concepts and their features (e.g., Quillian, 1967; Smith & Medin, 1981; Tyler, Moss, Durrant-Peatfield, & Levy, 2000; Rogers & McClelland, 2004). Our prediction is that the stronger the representation of the associated feature, the faster and more accurate should the concept tokening process be.




Figure 1: Schematic representation of the concept tokening process by vision and language; the cloud represents the featural hypothesis about the concept DOG, with circle size representing feature saliency.




Figure 2: Time-course of events for each trial in the picture-word priming congruency task. The words presented with the image were either a basic level label (e.g., *dog*), a high salient feature (e.g., *bark*), a low salient feature (e.g., *fur*), or a superordinate category label (e.g., *animal*).

This prediction is also informed by usage-based or distributional semantic models (DSMs) which place the burden of conceptual representation (viz., word meanings) on the patterns of lexical use. According to these models (e.g., Lund & Burgess, 1996; Landauer & Dumais, 1997; Johns, 2021b), the semantic/conceptual system is built by the strength of lexical co-occurrences in discourse, yielding representations in multidimensional space that capture the semantic relations between words (see, e.g., Jones, Kintsch,

& Mewhort, 2006). Crucially, DSM models take advantage of the *frequencies* of words across contexts, relying on large corpora, while emphasizing the role that these contexts play in building semantic representations and their relations.

We conducted generalized linear mixed-effects models to investigate the role of sample count-based measures (i.e., measures based on the norming study: sample frequency, cue validity, feature distinctiveness), and social usage count-based measures (i.e., measures based on corpus data: WF, CD, UCD, DCD) during the earliest moments of feature-to-concept tokening. Our goal was to understand what type of information might be accessed at the moment object and word recognition are attained simultaneously, that is, at the first 50-60 milliseconds of perceptual encoding (e.g., Potter, 2018).

Method

Picture-Word Priming Congruency Task

The behavioral data analyzed here was taken from a separate study employing a picture-word-priming congruency task (PWPC) task with 71 participants ($F = 52$; $M_{age} = 25$; $SD_{age} = 8$; Antal & de Almeida, 2022b). In this study, participants were presented with picture-word pairs to judge whether the items were related to each other. Pictures and target words were presented simultaneously with a 10 ms asynchrony accounting for their different recognition times, with objects presented for 50 or 190 ms and words for 60 or 200 ms. For each picture, one of four word probes was presented for congruency decision: the basic level category label (*dog*), a high-salient (most frequent: *bark*), a low-salient (less frequent: *fur*), or a superordinate category label (*animal*; see Figure 2). The rationale for using the PWPC task is that a decision to a picture-word pair relies on accessing the *content* of the word as it relates to the *content* of the depicted object.

In the present paper, we report analyses from accuracy responses to congruency decisions for the high- and low-salient features presented in the 50/60 ms condition. We focused on 50/60 ms condition because this timepoint allowed us to probe the nature of the information that is accessed during the early moments of object recognition. Crucially, investigating the earliest moments of concept tokening allows for a distinction between elucidating what kind of information arises *at the moment objects are recognized* (i.e., at the moment of concept “activation”), from information that arises from potential inferences *triggered* by the concept, but which may not be part of the content of the concept itself (viz., knowing that dogs bark; Fodor, 1983; de Almeida & Antal, 2021; de Almeida & Lepore, 2018).

Object concepts were divided into living and nonliving categories: half of the pictures represented images of living things (e.g., animals, vegetables, fruits), while the other half represented images of nonliving things (e.g., furniture, tools, clothing). This distinction was included for several reasons. First, living and nonliving superordinate categories allow for generalizations across major concept categories. Second, living things are assumed to share more features between

themselves, and encompass a larger number of features than nonliving things—potentially facilitating the tokening of a concept (Farah, McMullen, & Meyer, 1991; Moss, Tyler, & Taylor, 2007). Third, impairments to living and nonliving categories is one of the most well-documented kinds of double dissociations in cases such as Alzheimer’s (e.g., see de Almeida, Mobbayen, Antal, Kehayia, Nair & Schwartz, 2021) and brain lesions due to vascular accidents, traumas, or infections (e.g., see Mahon & Caramazza, 2009). As such, the pattern of dissociations obtained from patients with different aetiologies suggest that the functional architecture of the unimpaired conceptual system takes living and nonliving things as representing a major distinction between the representation of categories in the brain—thus, motivating a distinction in materials and analyses.

Features of Object Concepts: Norming Data

The lexical labels presented for each picture were determined through a norming study conducted with 100 participants ($F = 47$, $M_{age} = 39$; $SD_{age} = 12$; Antal & de Almeida, 2022a). In this study, participants were presented with all 260 line drawings from the Snodgrass and Vanderwart’s (1980) picture set and had to perform three tasks: (1) name the object depicted in the picture, as briefly and unambiguously as possible (e.g., dog); (2) categorize the object using a general term—targeting superordinate responses (e.g., animal), and (3) list three features (e.g., physical, sensory, functional properties) that came to mind, regarding the object that was presented to them. For all tasks, participants were instructed to respond using only one word for phase of the task, and to use the first word that came to mind (or the first three in the case of features). The semantic features and categorization norms allowed us to gather a total of 78,000 features.

High- and Low-Salient Features. Based on these norms, the high- and low-salient lexical labels were determined through a ranked weighted response system. Specifically, the feature that was listed first by a given participant received a score of 3, the second feature received a score of 2, and the last feature a score of 1. These scores were then multiplied by the number of participants responding a given feature in their ranked position. Finally, their products were summed across all ranked positions and divided by the total number of participants. This yielded the final naming agreement for a given target feature. For instance, in response to the picture ‘banana’, 87 individuals responded *yellow* as the first feature, 12 as the second feature, and 1 as the third feature. As such, 87 was multiplied by 3, 12 was multiplied by 2, and 1 was multiplied by 1. Their products were then summed (i.e., 286) and divided by the total number of participants (i.e., 100), for a naming agreement value of 2.86. Low-salient targets were determined by taking the feature corresponding to half of the naming agreement value of the high-salient feature. In cases where no feature precisely matched that naming agreement value, the feature with the closest lower value was taken to be the target. In cases of a tie between two features (i.e., those corresponding to precisely half of the naming agreement

value of the high-salient feature), the feature that was a constituent part of the object depicted in the picture was taken to be the target.

Cue Validity and Feature Distinctiveness. We also computed cue validity and distinctiveness values for each of the high- and low-salient features. *Cue validity* scores were devised by dividing the sum of the production frequency for a given feature for a particular picture (e.g., *fur* for the dog picture) by the total production frequency for that feature, across all 260 pictures (Rosch, 1978; Rosch & Mervis, 1975; Reed, 1972). *Distinctiveness* was defined as the inverse of the total number of pictures that a given features appears in (Tyler et al., 2013; Randall, Moss, Rodd, Greer, & Tyler, 2004; Devereux, Taylor, Randall, Geertzen, & Tyler, 2016).

Feature Subcategory. In addition to obtaining lexical labels, the norming study also allowed us to classify all lexical labels for the high and low salient features according to 10 properties that the labels express about the target object: visual (e.g., *red* for apple), taste (e.g., *sweet* for cake), sound (e.g., *purr* for cat), substance (e.g., *metal* for hammer), dimension (e.g., *large* for elephant), shape (e.g., *round* for ball), part-to-whole (e.g., *leg* for table), function (e.g., *warm* for glove), quality (e.g., *soft* for bed), and concept-association (e.g., *bread* for toaster).

Lexical Properties of Corpora Measures

We relied on four different social usage count measures of lexical strength in our models: (1) WF, (2) CD, (3) DCD, and (4) UCD. All were attained from a large corpus of Reddit comments analyzed by Johns (2021a). Word frequency is number of occurrences of a word across all comments in the Reddit corpus. CD is the number of comments a word occurred in (roughly analogous to a context size of a paragraph). The DCD count is the number of discourses (operationally defined here as a subreddit) that a word was used in (has a maximum value of 30,327, which is the total number of subreddits contained in the corpus), while UCD is the total number of users who used a word in their comments (a maximum value of 334,345, the total number of user in the corpus). Each variable used in the analysis was reduced with a natural logarithm, consistent with past research (Adelman & Brown, 2006; Jones, Johns, & Recchia, 2012).

Data Analyses

We first explored how the properties of lexical items, which label object features, may influence concept tokening. We investigated this question in two separate analyses: (1) one relying on the sample count-based measures and (2) a second relying on the social usage count-based measures. For the count-based measures, we hypothesized greater accuracy for object features that are stronger and more salient to the representation of the concept. In particular, we predicted that object features that are high in cue validity and distinctiveness would signal stronger category membership and would thus lead to more accurate concept tokening.

Moreover, for the social usage count-based measures, if lexical-semantic information is organized in large part due to the social environment (i.e., language use), we predicted that the UCD and DCD measures should lead to better concept tokening and should thus better predict accuracy to congruency decisions. We then conducted exploratory analyses on feature subcategories (e.g., visual, taste, sound, shape, substance, dimension, part-to-whole), aiming to understand the role played by these subcategories in concept tokening, as measured by accuracy on congruency decisions.

To investigate these questions, we analyzed participants' accuracy on the PWPC task with binomial generalized linear mixed-effects (GLME) regression models, using the *lme4* package (Bates, Maechler, & Bolker, 2013) for the R statistical programming environment (R. Dev Core Team, 2021). Our first model estimated the effect of sample count-based measures, by including living/nonliving, sample frequency, cue validity, and feature distinctiveness as fixed effects. For the social usage count-based measures, we fitted four different models, one for each of the corpora measures: (1) WF, (2) CD, (3) UCD, and (4) DCD. Each of these models included living/nonliving and the corpora measure of interest as fixed effects. Given the large variability in scaling, the corpora measures were log-transformed. As such, all of the social usage count-based models relied on log-transformed data. Our last model estimated the effect of feature subcategories, by entering living/nonliving and feature subcategory as fixed effects. All models included random intercepts for subjects and items, as justified by the likelihood tests. We derived *p*-values for all main effects and interactions using the Likelihood Ratio Test by comparing the full model with all our fixed effects of interest to a reduced model excluding the relevant term (Winter 2013, 2019). Planned comparisons were conducted using the *emmeans* package with Tukey's correction (Lenth et al., 2018). All figures were created using *ggplot2* (Wickham, 2016).

Results and Discussion

Prior to conducting analyses, participants' overall accuracy to congruency decisions were screened. All participants scored above chance (i.e., above 50%) and were thus kept for all analyses.

Sample Count-Based Measures

The full model was compared to a null model consisting of only random predictors and was found to provide a statistically significant better fit to the data, $\chi^2(4) = 83.92$, $p < 0.001$, $R^2 = 0.22$, 95% CI [0.00, 0.45]. There were also significant main effects of living/nonliving, sample frequency, and cue validity, as well as a living/nonliving by feature distinctiveness interaction (see Table 1). As predicted, participants were more accurate when presented with items from living categories. In particular, the odds of correct congruency decisions for living things were 1.61 times that of nonliving things. Interestingly, as features increased in cue validity, participants' accuracy significantly decreased. A similar effect was also found with feature distinctiveness, whereby participants' accuracy significantly




Figure 3: (A) Mean correct accuracy for congruency decisions between living and nonliving objects. Participants' accuracy as a function of log sample frequency (B), cue validity (C), and feature distinctiveness (D), across living and nonliving categories. Error bars correspond to 95% CI of group means.

Table 1. Generalized linear mixed effects regression for the sample count-based measures.

Predictors	β	SE β	z-value	OR	95% CI of OR	Null Comparison
Intercept	0.52	0.14	3.64	1.68	[1.27, 2.22]	
Living/Nonliving	0.47	0.14	3.52	1.61	[1.22, 2.11]	$\chi^2(1) = 11.10$, $p < 0.001$ *
Sample Frequency	0.01	0.01	5.60	1.01	[1.01, 1.02]	$\chi^2(1) = 31.45$, $p < 0.001$ *
Cue Validity	-0.74	0.29	-2.54	0.48	[0.27, 0.84]	$\chi^2(1) = 6.41$, $p = 0.01$ *
Distinctiveness	-0.73	0.52	-1.41	0.48	[0.17, 1.33]	$\chi^2(1) = 1.98$, $p = 0.16$
Living/Nonliving x Sample Frequency	-0.10	0.14	-0.69	0.91	[0.69, 1.20]	$\chi^2(1) = 0.47$, $p = 0.50$
Living/Nonliving x Cue Validity	-0.39	0.34	-1.13	0.68	[0.35, 1.33]	$\chi^2(1) = 1.26$, $p = 0.26$
Living/Nonliving x Distinctiveness	-1.30	0.62	-2.10	0.27	[0.08, 0.92]	$\chi^2(1) = 4.42$, $p = 0.04$ *

decreased as features increased in distinctiveness—an effect which was more pronounced for living things (see Figure 3).

In fact, these two measures seem to go in the same direction, as they are strongly correlated ($r = 0.65$). Although seemingly paradoxical, the living/nonliving by distinctiveness interaction effect has also been found by Randall et al. (2004). According to Randall et al., the disadvantage in processing distinctive properties of living things is due to their lack of correlation with other features (e.g., a lion's *mane* is weakly correlated with other features), in comparison to nonliving things, which express highly correlated distinctive features. We propose an alternative interpretation. As several studies have shown (Wu, Crouzet, Thorpe, & Fabre-Thorpe, 2014; Poncet & Fabre-Thorpe, 2014; Rogers & Patterson, 2007), under rapid categorization conditions (i.e., 60 ms or below, as in the present study), the conceptual system is attuned to general, superordinate information. The PWPC paradigm, with 60 ms presentation times, may constrain the quality and amount of perceptual information required for participants to judge the congruency between object and feature, which may depend on the processing of finer perceptual details. Therefore, presenting a specific feature (even a highly salient one) during the earliest moments of concept tokening may hinder accuracy, as one may require more time to analyze the constituent properties of the object concept. To wit, during the early stages, the conceptual system knows that a dog is an animal, but not that it barks.

Social Usage Count-Based Measures

Each of the corpora measure models provided a statistically significant better fit to the data than a null model consisting of only random predictors. Additionally, results revealed main effects for all four of the corpora measures, as well marginal interactions between living/nonliving and three of the corpora measures: WF, UCD, and DCD (see Table 2). As shown in Figure 4, while all of the corpora measures engendered greater accuracy for every unit increase in frequency, there is a noticeable advantage for the UCD and DCD variables, over CD and WF. In particular, the odds of correct congruency decision are markedly larger for the UCD and DCD variables (1.49 and 1.67, respectively), in comparison to those of WF and CD (both 1.17). Moreover, although there were no statistically significant differences




Figure 4: Participants' accuracy on congruency decisions as a function of word frequency (A), contextual diversity (B), user contextual diversity (C), and discourse contextual diversity (D), across living and nonliving categories.

between models for the individual variables, the UCD/DCD models explained a greater proportion of the variance in accuracy than the WF/CD models (R^2 : WF = 0.2179, CD = 0.2165, UCD = 0.2243, DCD = 0.2204). Our results suggest that concept tokening may be affected by properties from the social environment. That is, as one experiences a given lexical item across many discourse contexts, and across many individuals within their social environment, the properties of that lexical item at the discourse (DCD) and user level (UCD) may, over time, become associated within the lexical-semantic system as an organizing principle. Thus, words that are frequently used across individuals and discourses, in particular, lead to a facilitation in concept tokening.

Feature Subcategories

Our full model with fixed effects was found to provide a statistically significant better fit to the data than the null model consisting of only random predictors, $\chi^2(15) = 92.08$,

Table 2. Generalized linear mixed effects regression for the social usage count-based measures.

Predictors	β	SE β	z-value	OR	95% CI of OR	Null Comparison
Word Frequency (WF)	0.16	0.02	3.64	1.17	[1.12, 1.23]	$\chi^2(1) = 42.61, p < 0.001$ *
Contextual Diversity (CD)	0.16	0.03	6.34	1.17	[1.11, 1.23]	$\chi^2(1) = 41.23, p < 0.001$ *
User Contextual Diversity (UCD)	0.40	0.05	7.49	1.49	[1.34, 1.65]	$\chi^2(1) = 58.41, p < 0.001$ *
Discourse Contextual Diversity (DCD)	0.51	0.07	7.13	1.67	[1.45, 1.92]	$\chi^2(1) = 52.45, p < 0.001$ *
Living/Nonliving x WF	-0.09	0.05	-1.76	0.92	[0.83, 1.01]	$\chi^2(1) = 3.05, p = 0.08$.
Living/Nonliving x CD	-0.08	0.05	-1.67	0.92	[0.83, 1.01]	$\chi^2(1) = 2.75, p = 0.10$
Living/Nonliving x UCD	-0.20	0.11	-1.83	0.82	[0.66, 1.01]	$\chi^2(1) = 3.34, p = 0.07$.
Living/Nonliving x DCD	-0.28	0.15	-1.88	0.76	[0.57, 1.01]	$\chi^2(1) = 3.47, p = 0.06$.

$p < 0.001$, $R^2 = 0.23$, 95% CI [0.00, 1.00]. There were also significant main effects of living/nonliving ($\chi^2(1) = 10.12$, $p = 0.002$), feature subcategory ($\chi^2(14) = 82.63$, $p < 0.001$), as well as a living/nonliving by feature subcategory interaction ($\chi^2(9) = 19.25$, $p = 0.02$). There was a living things advantage (i.e., greater accuracy for living vs. nonliving) for items from the subcategory of *dimension* ($OR = 2.08$, $p = 0.02$), *quality* ($OR = 2.49$, $p = 0.006$), *substance* ($OR = 1.79$, $p = 0.04$), and *visual* ($OR = 3.45$, $p = 0.07$). Interestingly, a nonliving advantage was found for the subcategory of *sound* ($OR = 4.07$, $p = 0.01$). Together, these results suggest an advantage for features that express perceptual properties—those which are more distinct in living things. It is also interesting to note that the nonliving objects within the *sound* subcategory were predominantly comprised of musical instruments, which in the category-specific semantic deficits literature have been shown to pattern with living things (Warrington & Shallice, 1984; de Almeida et al., 2021, Zannino et al., 2007).

Conclusion

The goal of the present study was to investigate how the properties of lexical items, which label object features, affect concept tokening. We were interested in understanding the role of lexical strength of sample count-based measures (sample frequency, cue validity, feature distinctiveness) and social usage count-based measures (WF, CD, DCD, UCD) during the earliest moments of feature-to-concept tokening. Methodologically, our models relied on behavioral data from a PWPC task with short presentation times, allowing us to probe the properties of lexical items that label object features at a relatively ‘early’ point in the time course of conceptual processing. Beyond the sample-based measures, the inclusion of social usage-based measures allowed us to gain insight into the contribution of diverse socially oriented contextual measures of lexical items and how they may affect concept tokening. Our approach enabled us to understand the properties of the semantic system by relying on the combination of language usage data as well as behavioral data to tap into the earliest moments of concept tokening.

Regarding the sample count-based measures, our results showed that cue validity and feature distinctiveness were negative predictors of participant’s accuracy to congruency decisions. Results showed that participant’s accuracy decreased as features increased in cue validity and in distinctiveness. This effect was more pronounced for living things, than nonliving things. Our results also showed that participants were more accurate when presented with object concepts from living categories—in particular when objects were paired with lexical items labeling features from the subcategories of *dimension*, *quality*, *substance*, and *visual*. We also found a nonliving advantage for object features from the subcategory of *sound*—an effect which has been consistently found in individuals with category-specific semantic deficits. For the social usage count-based measures, while all corpora measures were positive predictors of participant’s accuracy to congruency decisions, results revealed a noticeable advantage for the UCD and DCD

measures, over WF and CD. Namely, the UCD and DCD measures accounted for a greater proportion of variance and engendered greater odds of correct responses in participant’s congruency decisions, in comparison to WF and CD.

Overall, our results seem to suggest that the conceptual system may be organized as a function of the *principle of likely need* similarly to the lexical system (see e.g., Adelman et al., 2006). That is, the features that occur frequently across objects, and whose labels are frequently used at the discourse and user level, have stronger representations and, consequently, lead to a facilitation in concept tokening. This may account for the UCD and DCD accuracy advantage over the WF and CD measures: if a feature label is used across all discourse types, and by many individuals, then that feature label should be more available within the lexicon, and thus should yield a faster interface with the conceptual system. A similar explanation can be made for the sample count-based measures. On the assumption that features which are highly shared (e.g., *legs* for living things) and which co-occur frequently (e.g., *eyes* and *mouth*) are repeatedly used in our social environment, the speed of concept tokening should reflect a graded increase that is a function of its usage, both at the discourse and user level. Therefore, given that the distinctive features of living things are not highly shared, nor do they co-occur frequently with other object features, they should engender slower interface with the conceptual system—an effect that was found in the present study (see also Tyler et al., 2007; Randall et al., 2004).

It remains to be determined, however, whether distinctive features, whose lexical labels co-occur frequently at the discourse or user level, lead to faster concept tokening. For instance, although *bark* and *udder* are distinctive features of dog and cow, respectively, the co-occurrence of “dog” and “bark” may be much greater than that of “cow” and “udder”, at the user and discourse level. In fact, a simple collocate search on COCA (Davies, 2009) revealed that “dog” and “bark” co-occur 416 times per million, in comparison to “cow” and “udder”, which occurred only 41 times. As such, it is plausible to expect distinctive feature labels that co-occur frequently across discourses and users to engender faster concept tokening than those that do not.

DSMs propose that lexical usage reflects semantic organization. These models postulate that “information about the meaning of *hammer* can be determined by observing the contexts in which it appears” (Jones et al., 2006). The methods we combine are important venues to explore the relationships between the lexical system and the concepts that words label. It is not clear whether our results reflect a parallelism—or even a causal relation—between the factors determining lexical organization, such as the principle of likely need, and the way concepts are organized. We plan to explore these relationships in future iterations of our models. How concept tokening is affected by the interaction of various properties, namely, distinctiveness, sharedness, correlational structure (i.e., co-occurrences), and usage pattern in the social world, at both the discourse and user level, is an important topic for future research.

Acknowledgments

We would like to thank three anonymous reviewers for their helpful comments on the first draft of this paper. This research was supported by fellowships from the Natural Sciences and Engineering Research Council of Canada (NSERC) and from the Social Sciences and Humanities Research Council of Canada (SSHRC).

References

- Adelman, J. S., & Brown, G. D. (2008). Modeling lexical decision: The form of frequency and diversity effects. *Psychological Review*, 115, 214–227.
- Antal, C., & de Almeida, R. G. (2022a, in preparation). Semantic properties and categorization norms for the 260 Snodgrass and Vanderwart's (1980) objects: a 40-year conceptual update to a classic set.
- Antal, C., & de Almeida, R. G. (2022b, in preparation). When you recognize a dog, you also know it is an animal: Evidence from rapid object categorization.
- Bates, D., Maechler, M., and Bolker, B. (2013). Lme4: linear mixed effects models using Eigen and Eigen. R package version 0.9999099–2. Available at: <http://CRAN.R-project.org/package=lme4>
- de Almeida, R. G., & Antal, C. (2021). How can semantics avoid the troubles with the analytic/synthetic distinction? In S. Löbner, T. Gamerschlag, T. Kalenscher, M. Schrenk, & H. Zeevat. (Eds.) *Concepts, Frames and Cascades in Semantics, Cognition and Ontology*. Springer.
- de Almeida, R. G., Mobayyen, F., Antal, C., Kehayia, E., Nair, V. P., & Schwartz, G. (2021). Category-specific verb-semantic deficits in Alzheimer's disease: Evidence from static and dynamic action naming. *Cognitive Neuropsychology*, 38(1), 1–26.
- Davies, M. (2009). The 385+ million word Corpus of Contemporary American English (1990–2008+): design, architecture, and linguistic insights. *International Journal of Corpus Linguistics*, 14(2), 159–190.
- de Almeida, R. G., & Lepore, E. (2018). Semantics for a module. In R. G. de Almeida & L. Gleitman (Eds.), *On modules, concepts, and language: Cognitive science at its core*. Oxford, UK: Oxford University Press.
- Devereux, B. J., Taylor, K. I., Randall, B., Geertzen, J., & Tyler, L. K. (2016). Feature statistics modulate the activation of meaning during spoken word processing. *Cognitive science*, 40(2), 325–350.
- Farah, M. J., McMullen, P. A., & Meyer, M. M. (1991). Can recognition of living things be selectively impaired? *Neuropsychologia*, 29(2), 185–193.
- Fodor, J. A. (1983). *The modularity of mind*. Cambridge, MA: MIT Press.
- Fodor, J. A., & Pylyshyn, Z. W. (2015). *Minds without meanings: An essay on the content of concepts*. Cambridge, MA: MIT Press.
- Johns, B. T. (2021a). Disentangling contextual diversity: Communicative need as a lexical organizer. *Psychological Review*, 128(3), 525–557.
- Johns, B. T. (2021b). Distributional social semantics: Inferring word meanings from communication patterns. *Cognitive Psychology*, 131, 1–20.
- Jones, M. N., Johns, B. T., & Recchia, G. (2012). The role of semantic diversity in lexical organization. *Canadian Journal of Experimental Psychology*, 66(2), 115–124.
- Jones, M. N., Kintsch, W., & Mewhort, D. J. (2006). High-dimensional semantic space accounts of priming. *Journal of memory and language*, 55(4), 534–552.
- Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological review*, 104(2), 211–240.
- Lenth, R., Singmann, H., Love, J., Buerkner, P., & Herve, M. (2018). Package “Emmeans”. R Package Version 4.0-3.
- Lund, K., & Burgess, C. (1996). Producing high-dimensional semantic spaces from lexical co-occurrence. *Behavior research methods, instruments, & computers*, 28(2), 203–208.
- Mahon, B. Z., & Caramazza, A. (2009). Concepts and categories: a cognitive neuropsychological perspective. *Annual review of psychology*, 60, 27–51.
- Moss, H. E., Tyler, L. K., & Taylor, K. I. (2007). Conceptual structure. In G. Gaskell (Ed.), *The Oxford handbook of psycholinguistics* (pp. 217–234). Oxford, UK: Oxford University Press.
- Poncet, M., & Fabre-Thorpe, M. (2014). Stimulus duration and diversity do not reverse the advantage for superordinate-level representations: the animal is seen before the bird. *European Journal of Neuroscience*, 39(9), 1508–1516.
- Potter, M. C. (2018). The immediacy of conceptual processing. In R. G. de Almeida & L. R. Gleitman (Eds.), *On concepts, modules, and language: Cognitive science at its core* (pp. 239–248). New York, NY: Oxford University Press.
- Quillian, M. R. (1967). Word concepts: A theory and simulation of some basic semantic capabilities. *Behavioral science*, 12(5), 410–430.
- R Core Team (2021). R: A Language and environment for statistical computing. *R Foundation for Statistical Computing*, Vienna, Austria. URL: <https://www.R-project.org/>.
- Randall, B., Moss, H. E., Rodd, J. M., Greer, M., & Tyler, L. K. (2004). Distinctiveness and correlation in conceptual structure: Behavioral and computational studies. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 30(2), 393–406.
- Reed, S. K. (1972). Pattern recognition and categorization. *Cognitive psychology*, 3(3), 382–407.
- Rogers, T. T., & McClelland, J. L. (2004). *Semantic cognition: A parallel distributed processing approach*. Cambridge, MA: MIT Press.
- Rogers, T. T., & Patterson, K. (2007). Object categorization: Reversals and explanations of the basic-level advantage. *Journal of Experimental Psychology: General*, 136(3), 451–469.

- Rosch, E. (1978). Principles of Categorization. In E. Rosch & B. Lloyd (Eds.), *Cognition and categorization* (pp. 27-48). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Rosch, E., & Mervis, C. B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive psychology*, 7(4), 573-605.
- Rosch, E., Mervis, C. B., Gray, W. D., Johnson, D. M., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive psychology*, 8(3), 382-439.
- Smith, E. E., & Medin, D. L. (1981). *Categories and concepts*. Cambridge, MA: Harvard University Press.
- Snodgrass, J. G., & Vanderwart, M. (1980). A standardized set of 260 pictures: norms for name agreement, image agreement, familiarity, and visual complexity. *Journal of experimental psychology: Human learning and memory*, 6(2), 174-215.
- Tyler, L. K., & Moss, H. E. (2001). Towards a distributed account of conceptual knowledge. *Trends in cognitive sciences*, 5(6), 244-252.
- Tyler, L. K., Chiu, S., Zhuang, J., Randall, B., Devereux, B. J., Wright, P., Clarke, A., & Taylor, K. I. (2013). Objects and categories: feature statistics and object processing in the ventral stream. *Journal of cognitive neuroscience*, 25(10), 1723-1735.
- Tyler, L. K., Moss, H. E., Durrant-Peatfield, M. R., & Levy, J. P. (2000). Conceptual structure and the structure of concepts: A distributed account of category-specific deficits. *Brain and language*, 75(2), 195-231.
- Warrington, E. K., & Shallice, T. (1984). Category specific semantic impairments. *Brain*, 107(3), 829-853.
- Wickham H (2016). *ggplot2: Elegant Graphics for Data Analysis*. New York: NY: Springer-Verlag New York.
- Winter, B. (2013). Linear models and linear mixed effects models in R with linguistic applications. *arXiv preprint arXiv:1308.5499*.
- Winter, B. (2019). *Statistics for linguists: An introduction using R*. Abington, UK: Routledge.
- Wu, C. T., Crouzet, S. M., Thorpe, S. J., & Fabre-Thorpe, M. (2015). At 120 msec you can spot the animal but you don't yet know it's a dog. *Journal of Cognitive Neuroscience*, 27(1), 141-149.
- Zannino, G. D., Perri, R., Caltagirone, C., & Carlesimo, G. A. (2007). Category-specific naming deficit in Alzheimer's disease: The effect of a display by domain interaction. *Neuropsychologia*, 45(8), 1832-1839.