Title
Partisan Representations: Partisan Differences in Semantic Representations and their Role in Attitude Judgements

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 40(0)

Authors
Halpern, David J
Rodriguez, Pedro L.K.

Publication Date
2018

Partisan Representations: Partisan Differences in Semantic Representations and their Role in Attitude Judgments

David J. Halpern (david.halpern@nyu.edu)
Department of Psychology, NYU
6 Washington Place
New York, NY 10003 USA

Pedro L. Rodríguez (pedro.rodriguez@nyu.edu)
Department of Politics, NYU
19 W. 4th St
New York, NY 10003 USA

Abstract

We outline a new method to explore differences in semantic representations between groups and apply it to a novel domain where we might expect to find such differences: *politics*. We hypothesize and find confirmatory evidence that individuals of opposite partisanship, as measured by party identification, have different semantic representations. We further evaluate whether differences in representations are predictive of attitude judgments as long suggested by constructivist theories of attitudes from social psychology. We find differences are indeed predictive of attitudes even after controlling for other strongly predictive covariates (party identification and ideology). In discussing our results we sketch out a broader theory of the role of semantic memory in attitude judgments.

Keywords: semantic memory; individual differences; attitudes; politics; concepts; modeling

Introduction

A growing body of research in the semantic memory literature has identified individual differences in semantic memory organization. What was thought to be largely static and "shared" (Ochsner, Kosslyn, Yee, Chrysikou, & Thompson-Schill, 2013) has been found to vary as function of expertise (Beilock, Lyons, Mattarella-Micke, Nusbaum, & Small, 2008), culture (Medin et al., 2006; Ji, Zhang, & Nisbett, 2004), native versus second language (Borodkin, Kenett, Faust, & Mashal, 2016), sensorimotor experience (Yee, 2017), development stage (Markman, 1994) and bodily differences (Thompson-Schill, Kan, & Oliver, 2006) among others. Individual differences in semantic memory are likely to have implications for downstream cognitive processes. We suggest that *making attitude judgments* is one such downstream process.

According to constructive models of attitudes from social psychology, making attitude judgments involves sampling (consciously or subconsciously) a limited number of relevant associated concepts (or associations) from memory and computing a summary of the valence of the retrieved associations (Tourangeau, 1992; Zaller & Feldman, 1992; Lord & Lepper, 1999). Although memory retrieval is central to these models, it has never been the direct object of study, instead its role has been limited to providing a conceptual framework to empirical studies of experiment context (Judd, Drake, Downing, & Krosnick, 1991; Tourangeau & Rasinski, 1988) with some exceptions (Bhatia, 2017; Lenton, Sedikides, & Bruder, 2009). We suggest that semantic memory is likely the source of the considerations and therefore we should expect that differences in semantic memory retrieval will predict differences in the resulting summary and expressed attitudes. In particular, if constructive attitude models are correct, the valences associated with the retrieved associations should explain much of the variance in expressed attitudes.

In this paper we identify differences in representations in a novel domain that we argue is well suited to explore the effect of these differences on attitude judgments: *politics*. In doing so we also showcase a new method to systematically explore differences in representations between groups by estimating semantic representations directly from semantic fluency data.

For the purposes of this paper, we define a subject's semantic representation as an object that constrains the likelihood of concepts or considerations being retrieved from memory. While we focus on semantic network representations in this paper, we leave generalizing our results to alternate models, such as a semantic space model (Landauer & Dumais, 1997) or topic model (Griffiths, Steyvers, & Tenenbaum, 2007), to future work. In the paper, we often refer to a representation for a particular concept c by which we mean the particular region of the semantic representation in the neighborhood of c that is typically retrieved in a task.

Why Politics?

We test two hypotheses about the relationship between semantic memory and (political) attitudes:

1. Individuals of opposite partisanship (here defined by party identification) have different semantic representations for politically charged concepts.

2. An *individual's* semantic representation for a particular political concept will be predictive of *that individual's* expressed attitude judgments on topics related to that concept.

There are good reasons to expect individual differences in semantic representations as function of partisanship. Political scientists have identified consistent differences in the vocabulary used by political elites (Gentzkow, Shapiro, & Taddy, 2016) and media organizations (Morris, 2007) as a function of political affiliation. Moreover, voters' media consumption habits have also been found to show a preference for media outlets perceived to be aligned with currently held political views (Mitchell, Gottfried, Kiley, & Matsa, 2014). Together these findings suggest two individuals of opposite partisanship are likely to have very different linguistic experiences. A fundamental prediction of linguistic based theories of concept acquisition is that differences in linguistic experience will produce different representations (Vigliocco, Meteyard, Andrews, & Kousta, 2009; Steyvers, Griffiths, & Dennis, 2006).

Differences in linguistic experience need not be the only source of representational differences in political concepts. Recent work highlights the role of emotions or affect as another type of experiential information relevant in forming semantic representations (Vigliocco et al., 2009; Ponari, Norbury, & Vigliocco, 2017), particularly for abstract concepts such as those we are likely to find in politics (e.g. "freedom", "peace" etc.). To the extent that individuals experience different emotions when partaking in political activities or encountering political content (Westbury, Keith, Briesemeister, Hofmann, & Jacobs, 2015), we should again expect differences in representations to emerge and, more to the point, differences that are likely to be highly relevant for attitude judgments.

Data

To evaluate these hypotheses, we need to estimate the semantic representations of political concepts for various partisan groups. In the semantic memory literature, semantic spaces are often estimated from large text corpora (Lund & Burgess, 1996; Landauer & Dumais, 1997) or a large set of word associations (Austerweil, Abbott, & Griffiths, 2012; De Deyne, Navarro, Perfors, & Storms, 2016). These methods are undesirable for our setting since we want to estimate the semantic representations of various sub-populations (members of political parties), something that would be difficult to do with large text corpora since it is unclear how to select a corpus for each sub-population and topic of interest. In addition, we are interested in topics where we have very weak priors on the extent of the semantic space (relative to more common topics like "fruits") so collecting word associations would require extensive and expensive piloting.

Instead, we build on a literature that estimates semantic representations from the semantic fluency task whereby participants are provided a category label as a cue (e.g. animals, food) and are asked to list as many examples of that category as they can think of within a given time limit and without repetition (Bousfield & Sedgewick, 1944). The semantic fluency task is ideal for our purposes for several reasons: First, in contrast to corpora, semantic fluency lists can be targeted to specific sub-populations of interest, better capturing group idiosyncrasies. Second, it allows us to quickly collect lots of data per subject that is relevant for the given category without the need for priors on which words to use to explore that category. Third, in addition to data on associations, it gives us data on the semantic memory search process which we hypothesize is relevant to predicting attitudes. Fourth, it has been shown to produce better models of semantic representations than single word associations (De Deyne, Navarro, & Storms, 2013).

We collected semantic fluency data from 1056 MTurk subjects. As cues we selected words that are politically relevant: welfare, government, American values, Republican and Democrat. For each cue, subjects were required to respond with associated words without repetitions.[1] Subjects also answered a series of demographic and political attitudes questions, including party identification and ideology.[2] We apply some basic preprocessing to the lists including spelling check, lower casing and singularizing basic plurals (e.g. "patriots" becomes "patriot"). Table 1 provides summary statistics of the resulting lists segregating by party identification.

Method

To estimate partisan differences in representations, we propose a new method that can be used to identify differences in group representations in general. We first estimate separate representations for Democrats and Republicans from their respective semantic fluency lists. Next, we compare the likelihoods of a set of heldout lists under each estimated representation. If there are partisan differences and individuals of the same partisanship overlap more in their representations than individuals of opposing partisanship, then the likelihood of Republican (Democrat) heldout lists should be larger under the Republican (Democrat) representation (within-party) than under the Democrat (Republican) representation (across-party).

We here assume that a semantic representation is a

[1] Although not a typical category fluency task, the task can be framed as one with the category defined as "words you associate with the cue".

[2] Possible answers to party identification included: "Republican", "Democrat", "Independent", "no preference" and "Other party (please specify)". Ideology was measured on a seven-point Likert scale from "Extremely liberal" to "Extremely conservative" with the option of choosing "Haven't thought much about this".

Table 1: summary statistics for concepts lists by party identification after pre-processing.

	Welfare		Government		American Values		Democrat		Republican	
	D	R	D	R	D	R	D	R	D	R
# of unique tokens	1291	912	1552	1085	1507	966	1347	1116	1513	992
Prop. overlap	0.277	0.367	0.305	0.370	0.313	0.339	0.281	0.416	0.303	0.352
Mean list length	14.105	12.848	15.194	14.561	14.354	14.251	14.266	13.696	14.677	14.133
	(4.499)	(4.665)	(4.384)	(4.322)	(4.432)	(4.558)	(4.294)	(4.544)	(4.254)	(4.352)

Notes: D = Democrats and R = Republicans. Prop. overlap corresponds to the ratio of unique tokens to total tokens listed. Numbers in parentheses denote standard deviations of mean list length.

network that is parametererized by an initial probability vector π which contains the probabilities of jumping from the cue word (e.g. "animal") to a given node (e.g. "dog") and a transition matrix $\mathbf{P}$ where each element of the matrix P_{ij} represents the probability of transitioning from word i to word j (e.g. from "dog" to "cat") in one step on the network. For a given π and $\mathbf{P}$, we can compute the likelihood of each list in our dataset and use maximum likelihood or Bayesian inference to infer the parameters of our semantic network. In the past, estimating representations in this way was not possible because the requirement that no word be repeated makes the likelihood of a true generative model non-trivial to compute. Previous models were either non-generative (e.g. Goñi et al. (2011)) and could not give likelihoods or were biased in their estimation process (Millsap & Meredith, 1987). Only recently has a generative model been proposed which could give likelihoods of producing semantic fluency lists under a set of estimated parameters (which determine the semantic representation). Jun, Zhu, Rogers, and Yang (2015) show that by assuming a particular model for the search process, they can estimate the semantic representation of a group that predicts new lists better than previous biased methods. Building on Austerweil et al. (2012), Jun et al. propose a model, called INVITE, whereby retrieval consists of a random walk through the semantic network with words being added to the semantic fluency list every time it reaches a new node. However, due to the constraints of the task, a word that has already been said cannot be repeated so if the random walk reaches a node that corresponds to a repeated word, no word is emitted.

By using the same generative model for both groups, Democrats and Republicans, we are assuming that there are no systematic differences in the search algorithm employed to retrieve associations. We argue that the search algorithm is likely to be a more fundamental cognitive process independent of individual differences in party identification. In Halpern and Rodriguez (2018) we tested this assumption by comparing the performance of several different models estimated separately on the two groups. The ranking of models according to the log-likelihood of held-out lists was the same for both groups, lending support to our assumption. [3]

Individual Differences

We divide up our data into 10 folds, stratifying on party identity. Estimation of the networks is easier and more reliable if it is limited to words that were included in several subjects' lists. Given the spread of words that subjects used, we restrict our estimation to the top 30 tokens said for each topic. We estimated a maximum likelihood "population semantic network" for self-identifying Democrats and Republicans (using LBFGS in rStan (Carpenter et al., 2016)) on a training set of 9 of the folds and then evaluated the log-likelihood of the heldout fold under each of these two semantic networks. Figure 2 plots an example of an estimated Democrat semantic representation for the concept *Republican*. Across all ten heldout folds, for all concepts and both parties we find that the within-party log-likelihood is significantly higher than across-party log-likelihood. As a measure of how well our model is able to differentiate parties, we can treat our model as a Bayesian classifier and assign the party with the higher log-likelihood to each list. Figure 1 plots the average accuracy of this classifier by concept. In all cases, the classifier is able to perform significantly better than chance [4]. The accuracy score in this case has a theoretically substantive interpretation: the larger the representational differences between groups, the easier it is for a classifier to distinguish between a Republican and Democrat resulting in a larger accuracy score, for political scientists this can be understood as a measure of "polarization" (Peterson & Spirling, n.d.). To further benchmark our results we applied our method to estimate semantic representations by gender rather than by party.[5] Figure 1 also plots the accuracy scores by gender for each concept. Except for the concept "Republican", our results suggest no significant gender differences in representations for our set of cues. Overall finding is evidence that Democrats and Republicans do indeed strongly differ in their represen-

[3]In this model comparison, we found that INVITE yields better results than many simpler models (including a simple bag-of-words). Since it provides a good description of semantic memory retrieval and has been shown to have nice statistical properties (Jun et al., 2015), we focus on INVITE for our analyses here.

[4]Since our sample is stratified by party, chance is an accuracy score of 0.5.

[5]Gender has also been previously identified as a potential source of differences in semantic representations (Capitani, Laiacona, & Barbarotto, 1999).

tations for these concepts.

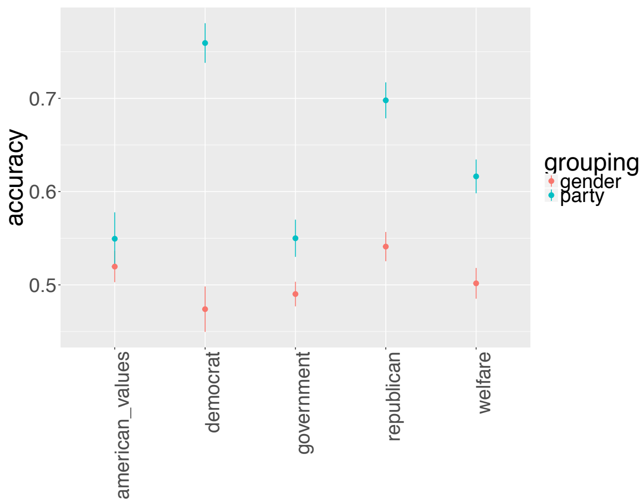

Figure 1: Accuracy of the model in discriminating between Democrat and Republican heldout subjects

Individual Differences and Attitudes

To explore whether retrieved semantic associations are predictive of attitude judgments we also collected data on general attitudes toward the government's role in providing services (related to the concept *government*) and its role in guaranteeing a minimum standard of living (related to the concept *welfare*). Both attitude questions were on a seven-point Likert scale and were recoded to range from -3 (extremely liberal position) to 3 (extremely conservative position).[6] We hypothesized the difference in the log-likelihoods of an individual's category fluency data under the Republican (LL_R) and Democrat (LL_D) models, a quantity we term *concept partisanship*, should be predictive of that individual's attitude judgments (using the representations for *welfare* for the question on welfare and *government* for the question on government services). The more negative (positive) the concept partisanship for subject i for concept c, the better that subject's fluency list approximates the Democrat's (Republican's) estimated representation. Table 2 reports our results of including concept partisanship as a regressor of expressed attitudes. Concept partisanship is significant even after controlling for party affiliation and ideology suggesting our representations are capturing more than group affiliation.

According to constructive models of attitudes, when responding to a survey question on attitudes individual's sample from memory, compute a statistic (e.g. an average) of the valences of the sampled information and respond accordingly. Building on this intuition we next asked how much predictive leverage can we get from simply using the average valence of the retrieved lists to predict attitude judgments. This requires we first attach a valence to the retrieved words which we do using a set of 13,915 valence norms from Warriner, Kuperman, and Brysbaert (2013). These valence norms range from a low of 1 ("unhappy") to a high of 9 ("happy") and subjects are instructed to respond how a word makes them feel. We emphasize this is an imperfect measure of valence to the extent that the valence of a word may change as a function of context and party affiliation yet it provides for an acceptable first approximation. Our results confirm that average valence of the retrieved lists is a significant predictor of expressed attitudes consistent with constructive models of attitudes (Table 2).

Table 2: Subjects' attitudes towards welfare as a function of partisanship of welfare concept representation and average valence of the subjects' retrieved lists

	Dependent variable:			
	Welfare Attitude			
	(1)	(2)	(3)	(4)
Concept Partisanship	0.335***	0.097***		
	(0.023)	(0.028)		
Average Valence			−0.662***	−0.307***
			(0.069)	(0.058)
Party (Republican = 1)		0.242		0.522**
		(0.283)		(0.261)
Ideology		0.487***		0.451***
		(0.063)		(0.063)
Constant	−0.123*	−0.094	2.952***	1.233***
	(0.070)	(0.152)	(0.347)	(0.319)
Observations	575	573	591	589
R^2	0.270	0.453	0.136	0.449
Adjusted R^2	0.269	0.450	0.134	0.446

Notes: Ideology ranges from -3 (extremely liberal) to 3 (extremely conservative).

Previous studies using semantic fluency tasks have observed that subjects produce items in bursts of semantically related words (Troyer, Moscovitch, & Winocur, 1997), consistent with semantic memory being organized in clusters of semantically related concepts. Given we found average valence of retrieved lists to be predictive of attitudes, we wondered whether valence serves as an organizing principle of semantic memory alongside semantic similarity (Osgood, Suci, & Tannenbaum, 1978; Westbury et al., 2015). One way of testing this hypothesis is to first assess whether clusters are present in our estimated representations and, given clusters are present, whether nodes within clusters tend to align according to valence. To evaluate the presence of clusters in our estimated representations we applied the *Walktrap* algorithm (Steinhaeuser & Chawla, 2010). Intuitively this algorithm identifies as clusters the densely connected regions of a graph in which simulated random walks tend to get "trapped".[7] Figure 2 plots the estimated Democrat semantic representation for the concept *Republican* with different colors representing different clusters. *Walktrap* algorithm identifies three distinct clusters. We draw the

[6] Both attitude questions were taken from the American National Elections Studies (ANES) Survey.

[7] The algorithm works best with small step sizes. We limit our random walk to 3 steps.

Table 3: Mean valence by cluster in Figure 2

	yellow	blue	green
mean	5.62	5.77	2.68
valence	(1.0250)	(1.390)	(0.567)

reader's attention to the highly negatively valenced *green* cluster vis-a-vis the other relatively more neutral clusters. Using the same valence norms we used in the regressions above, we can estimate mean valence by cluster (see Table 3). The green cluster (consisting of the words "corrupt", "greedy", "ignorant", "liar", "racist", "selfish" and "uncaring") is significantly more negatively valenced than the other two clusters. We see this as suggestive evidence of valence serving as an important organizing dimension of semantic memory, a result meriting further research.

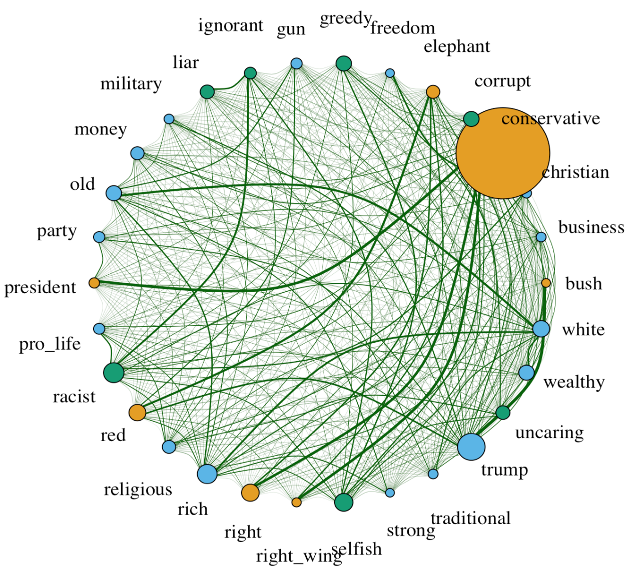

Figure 2: Democrat network for the concept "Republican." Clusters of concepts as estimated by the Walktrap algorithm are indicated by color

Discussion

We have outlined a method to explore differences in semantic representations between groups and applied it to a novel domain: politics. We hypothesized that individuals of opposite partisanship have different semantic representations for political relevant concepts. In our data, we find evidence of differences across several political concepts although the magnitude of the difference is found to vary by concept, with concepts related to self-identity (Democrat and Republican) showing the largest differences. We also hypothesized that an individual's semantic representation of a politically relevant concept is predictive of that individual's attitudes toward topics related with that concept. Again, we find strong confirmatory evidence of this hypothesis. Finally we also found evidence consistent with valence playing an important role, alongside semantic similarity, in the organization of semantic memory.

We began by arguing that partisan differences in representations are likely to have emerged as a result of differences in the linguistic and emotional experiences of Democrats and Republicans. We now proceed to sketch out a more general theory of the relationship between semantic memory and attitudes. We hypothesize, that there might be a computational reason for these differences that further constrains how representations develop and change. The organization of semantic memory is thought to be optimized for making efficient and accurate knowledge-based inferences and predictions (e.g. top-down perception (Biederman, Kubovy, & Pomerantz, 1981) and linguistic prediction (Steyvers et al., 2006). This is consistent with the fact that semantic memory has been found to be organized according to similarity in sensorimotor experiential data and language-based distributional data (Andrews, Vigliocco, & Vinson, 2009). However, many studies have suggested valence as another important dimension of semantic organization (Osgood et al., 1978; Westbury et al., 2015), potentially resulting from co-occurrence statistics of affective experience (Vigliocco et al., 2009).The fact that many of our most discriminating tokens are valenced and that similarly valenced nodes seem to cluster together is consistent this theory. This begs the questions: what use is valence as an organizing principle? We hypothesize that semantic memory is also optimized for efficient and consistent evaluative judgments under limited resources. If evaluative judgments do indeed follow a sampling like process then it makes sense for valence to play an organizing role lest individuals produce an endless stream of conflicting evaluations. We see this as a line research meriting greater attention and believe politics as a domain is ideally suited to this task. More generally we hope the method outlined above provides a basic framework to begin to quantitatively explore the relationship between semantic memory and attitudes and that our promising results serve to highlight the potential returns to cognitive science of branching into less traditional domains.

Acknowledgments

This research was generously supported by the George Downs Prize grants as well as an NSF grant BCS-1255538 and a John S. McDonnell Foundation Scholar Award to Todd M. Gureckis.

References

Andrews, M., Vigliocco, G., & Vinson, D. (2009). Integrating experiential and distributional data to learn semantic representations. *Psychological review*, *116*(3), 463.

Austerweil, J. L., Abbott, J. T., & Griffiths, T. L. (2012). Human memory search as a random walk in a semantic network. In *Advances in neural information processing systems* (pp. 3041–3049).

Beilock, S. L., Lyons, I. M., Mattarella-Micke, A., Nusbaum, H. C., & Small, S. L. (2008). Sports experience changes the neural processing of action language. *Proceedings of the National Academy of Sciences*, *105*(36), 13269–13273.

Bhatia, S. (2017). The semantic representation of prejudice and stereotypes. *Cognition*, *164*, 46–60.

Biederman, I., Kubovy, M., & Pomerantz, J. (1981). Perceptual organization. *On the semantics of a glance at a scene*, 213–263.

Borodkin, K., Kenett, Y. N., Faust, M., & Mashal, N. (2016). When pumpkin is closer to onion than to squash: The structure of the second language lexicon. *Cognition*, *156*, 60–70.

Bousfield, W. A., & Sedgewick, C. H. W. (1944). An analysis of sequences of restricted associative responses. *The Journal of General Psychology*, *30*(2), 149–165.

Capitani, E., Laiacona, M., & Barbarotto, R. (1999). Gender affects word retrieval of certain categories in semantic fluency tasks. *Cortex*, *35*(2), 273–278.

Carpenter, B., Gelman, A., Hoffman, M., Lee, D., Goodrich, B., Betancourt, M., ... others (2016). Stan: A probabilistic programming language. *Journal of Statistical Software*, *20*(2), 1–37.

De Deyne, S., Navarro, D. J., Perfors, A., & Storms, G. (2016). Structure at every scale: A semantic network account of the similarities between unrelated concepts. *Journal of Experimental Psychology: General*, *145*(9), 1228.

De Deyne, S., Navarro, D. J., & Storms, G. (2013). Better explanations of lexical and semantic cognition using networks derived from continued rather than single-word associations. *Behavior research methods*, *45*(2), 480–498.

Gentzkow, M., Shapiro, J. M., & Taddy, M. (2016). *Measuring polarization in high-dimensional data: Method and application to congressional speech* (Tech. Rep.). National Bureau of Economic Research.

Goñi, J., Arrondo, G., Sepulcre, J., Martincorena, I., de Mendizábal, N. V., Corominas-Murtra, B., ... Wall, D. P. (2011). The semantic organization of the animal category: evidence from semantic verbal fluency and network theory. *Cognitive processing*, *12*(2), 183–196.

Griffiths, T. L., Steyvers, M., & Tenenbaum, J. B. (2007). Topics in semantic representation. *Psychological review*, *114*(2), 211.

Halpern, D., & Rodriguez, P. (2018). Comparing models of semantic search in concrete and abstract categories. PsyArXiv.

Ji, L.-J., Zhang, Z., & Nisbett, R. E. (2004). Is it culture or is it language? examination of language effects in cross-cultural research on categorization. *Journal of personality and social psychology*, *87*(1), 57.

Judd, C. M., Drake, R. A., Downing, J. W., & Krosnick, J. A. (1991). Some dynamic properties of attitude structures: Context-induced response facilitation and polarization. *Journal of Personality and Social Psychology*, *60*(2), 193.

Jun, K.-S., Zhu, X., Rogers, T. T., & Yang, Z. (2015). Human memory search as initial-visit emitting random walk. In *Advances in neural information processing systems* (pp. 1072–1080).

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.

Lenton, A. P., Sedikides, C., & Bruder, M. (2009). A latent semantic analysis of gender stereotype-consistency and narrowness in american english. *Sex Roles*, *60*(3), 269–278.

Lord, C. G., & Lepper, M. R. (1999). Attitude representation theory. In *Advances in experimental social psychology* (Vol. 31, pp. 265–343). Elsevier.

Lund, K., & Burgess, C. (1996). Producing high-dimensional semantic spaces from lexical co-occurrence. *Behavior research methods, instruments, & computers*, *28*(2), 203–208.

Markman, E. M. (1994). Constraints on word meaning in early language acquisition. *Lingua*, *92*, 199–227.

Medin, D. L., Ross, N. O., Atran, S., Cox, D., Coley, J., Proffitt, J. B., & Blok, S. (2006). Folkbiology of freshwater fish. *Cognition*, *99*(3), 237–273.

Millsap, R. E., & Meredith, W. (1987). Structure in semantic memory: A probabilistic approach using a continuous response task. *Psychometrika*, *52*(1), 19–41.

Mitchell, A., Gottfried, J., Kiley, J., & Matsa, K. (2014). *Political polarization & media habits. pew research center.* Retrieved 11/5/2014 from http://www. journalism. org/2014/10/21/political-polarization-media-habits.

Morris, J. S. (2007). Slanted objectivity? perceived media bias, cable news exposure, and political attitudes. *Social Science Quarterly*, *88*(3), 707–728.

Ochsner, K. N., Kosslyn, S., Yee, E., Chrysikou, E. G., & Thompson-Schill, S. L. (2013). *Semantic memory. in the oxford handbook of cognitive neuroscience, volume 1: Core topics.* Oxford University Press.

Osgood, C. E., Suci, G. J., & Tannenbaum, P. H. (1978). The measurement of meaning. 1957. *Urbana: University of Illinois Press*.

Peterson, A., & Spirling, A. (n.d.). Classification accuracy as a substantive quantity of interest: Measuring polarization in westminster systems.

Ponari, M., Norbury, C. F., & Vigliocco, G. (2017). Acquisition of abstract concepts is influenced by emotional valence. *Developmental science*.

Steinhaeuser, K., & Chawla, N. V. (2010). Identifying and evaluating community structure in complex networks. *Pattern Recognition Letters*, *31*(5), 413–421.

Steyvers, M., Griffiths, T. L., & Dennis, S. (2006). Probabilistic inference in human semantic memory. *Trends in cognitive sciences*, *10*(7), 327–334.

Thompson-Schill, S. L., Kan, I. P., & Oliver, R. T. (2006). Functional neuroimaging of semantic memory. *Handbook of functional neuroimaging of cognition*, *2*, 149–190.

Tourangeau, R. (1992). Context effects on responses to attitude questions: Attitudes as memory structures. In *Context effects in social and psychological research* (pp. 35–47). Springer.

Tourangeau, R., & Rasinski, K. A. (1988). Cognitive processes underlying context effects in attitude measurement. *Psychological bulletin*, *103*(3), 299.

Troyer, A. K., Moscovitch, M., & Winocur, G. (1997). Clustering and switching as two components of verbal fluency: evidence from younger and older healthy adults. *neuropsychology*, *11*(1), 138.

Vigliocco, G., Meteyard, L., Andrews, M., & Kousta, S. (2009). Toward a theory of semantic representation. *Language and Cognition*, *1*(2), 219–247.

Warriner, A. B., Kuperman, V., & Brysbaert, M. (2013). Norms of valence, arousal, and dominance for 13,915 english lemmas. *Behavior research methods*, *45*(4), 1191–1207.

Westbury, C., Keith, J., Briesemeister, B. B., Hofmann, M. J., & Jacobs, A. M. (2015). Avoid violence, rioting, and outrage; approach celebration, delight, and strength: Using large text corpora to compute valence, arousal, and the basic emotions. *The Quarterly Journal of Experimental Psychology*, *68*(8), 1599–1622.

Yee, E. (2017). Fluid semantics: Semantic knowledge is experience-based and dynamic. *The Speech Processing Lexicon: Neurocognitive and Behavioural Approaches*, *22*, 236.

Zaller, J., & Feldman, S. (1992). A simple theory of the survey response: Answering questions versus revealing preferences. *American journal of political science*, 579–616.