Computational *Context* Analysis:
A High-Dimensional Cognitive Modeling Complement to Content Analysis

Chad Murphy
Doctoral Candidate
Department of Political Science

Martin Johnson*
Associate Professor
Department of Political Science

Curt Burgess
Professor
Department of Psychology
University of California, Riverside
Riverside, CA 92521

Abstract

We present a high-dimensional theoretical model of cognitive representation, focused on concept acquisition. For students of large-corpus text analysis and text annotation, this model will be particularly interesting due to the computational instantiation it implies, which allows cognitive, political, and social researchers to systematically measure and study meaning in text. We develop replicable measures of cognitive belief structures, the semantic similarity of words, ambiguity, among other attributes of text-based expressions. The approach we describe complements the more traditional aims of content analysis, such as categorization. We discuss several applications of high-dimensional cognitive modeling to political science research questions, with a particular focus on priming racial stereotypes and the assessment of ideological positions in text.

*Corresponding author: Martin Johnson, martin.johnson@ucr.edu. The authors thank Shaun Bowler for useful conversations and Tom Hayes and Hannah Postiff for help with data collection and preparation. All errors remain our own.

Computational *Context* Analysis:
A High-Dimensional Cognitive Modeling Complement to Content Analysis

We present a method for statistically analyzing large text samples rooted in a rich, parsimonious, cognitive theory of human concept acquisition and representation. Our approach complements other techniques informed by sophisticated statistical models, but representing thinner theories of cognition. Second, the model we present embraces meaning from a bottom-up, connotative perspective, rather than a top-down, denotative approach that assigns words to meanings based on dictionary conventions. Grounded in recent developments in concept representation in cognitive science, our approach emphasizes the role of context in language use. We introduce the Hyperspace Analogue to Language (HAL) model developed in cognitive science and discuss its 15-year track record of use in cognitive psychology, psycholinguistics, and natural language processing research.

Our theoretical model informs a computational approach to analyzing text that approximates the process of human learning from exposure to language. It embodies the expectations that a person understands the meaning of concepts, develops beliefs, and experiences a mental organization of beliefs and attitudes as a function of a person's experience consuming information. The ability to systematically, computationally, and reliably characterize the meaning of concepts and associations among objects in memory gives us access to a broad array of research questions.

In this paper, we discuss the theoretical assumptions of the HAL model, as well as the computational operationalization it informs. We compare and contrast HAL to several other approaches to text analysis, including a related model, Latent Semantic Analysis (Landauer & Dumais, 1997), which has served as inspiration for some political science research (Simon &

Xenos, 2004). We also present a pair of examples demonstrating our main point, that HAL mimics the human semantic system and reproduces other measures useful to political scientists using a large text sample as its principal input. These examples focus on racial stereotyping and the placement of politicians in ideological space. Finally, we discuss several other potential applications of the model.

Approaches to Text Analysis in Political Science

Political scientists have developed many ways of analyzing text. For example, hand-coded content analysis (Budge, Robertson, & Hearl, 1987; Hill, Hanna, & Shafqat, 1997), computerized approaches like WORDSCORES (Laver, Benoit, & Garry, 2003), and ideal point estimation (Monroe & Maeda, 2004) are quite useful for locating candidates in policy space. Dictionary-based approaches such the General Inquirer (Stone, Dunphy, Smith, & Ogilvie, 1966), DICTION (Hart, 2000), and a similar approach from psychology, the Linguistic Inquiry and Word Count (Pennebaker & Graybeal, 2001), categorize documents, and can identify themes and elements of personality, as can other political science approaches strongly influenced by machine-learning and Natural Language Processing that incorporate supervised or unsupervised learning (Hopkins & King, 2007; Purpura & Hilliard, 2006).

Many computerized approaches to content analysis base inferences on the frequency with which individual words appear in text. For example, Hart (2000) significantly advanced the study of campaign rhetoric using his DICTION software, which allows investigators to base inferences about any political actor's rhetorical style on the presence and frequency of the words they use. One of Hart's measures, CERTAINTY, is based on several sub-dictionaries, such as TENACITY, which sums up a person's use of being verbs (*is*, *am*, *will*) and related usages, such as contractions including these verbs, because "they connote confidence and totality" (Hart, 2000, p. 246). The validity of this measure hinges – in part – on the consistent use of being verbs to express certainty. Strategic users of language recognize that even the meaning of being verbs such as *is*, *was,* and *will* can vary by speaker and context (Starr, 1998) Most political scientists intentionally ignore the context in which words appear (e.g., Hart, 2000; Laver et al., 2003; Quinn, Monroe, Colaresi, Crespin, & Radev, 2007).

Further, most contemporary approaches to content analysis articulate few explicit assumptions about cognitive process, if any. Several content analytic approaches characterize word selection as a probabilistic process, allowing us to identify the probability that use of a given word suggests the likelihood we are reading a given reference text, in the case of WORDSCORES (Laver et al., 2003), or the probability the use of a given word suggests the identity of a particular speaker in the case of the IRT approaches (Monroe & Maeda, 2004; Slapin & Proksch, 2007). However, most of the underlying cognitive assumptions are unstated. Implicitly and explicitly, these approaches rely on the "bag of words" assumption, that text samples are composed of words expressed independently of one another (Quinn et al., 2007). Words are the behavior of interest, rather than the thought processes behind the selection of words or the contexts in which they are used. The ALCESTE content analysis approach comes somewhat closer to ours – it is sensitive to associations among concepts in rhetoric (Schonhardt-Bailey, 2005). LSA is a much more similar approach to HAL – it is also rooted in a high-dimensional model of concept acquisition.

An issue of *Discourse Processes* (Issues 2-3, 1998) was dedicated to research using HAL and LSA and covered a variety of both theoretical and applied problems. An important aspect of high-dimensional memory and language models such as HAL and LSA is that they learn with an unsupervised, purely inductive, algorithm. They share this characteristic with several approaches. For example, when given the same input, HAL and a simple recurrent network (a type of neural network) produce the same results due to the use of context in learning (Burgess & Lund, 2000). Perhaps the most important difference is that HAL is more appropriate for the study of word-to-word relationships, while LSA is useful to study the relationship between words and other textual units (e.g., paragraphs, sentences, etc.) Most of the objects of study that we are interested in will involve word-to-word relationships in the expressions of a single speaker or type of speaker, rather than relationships between words and other textual units and types of research problems more appropriate for LSA. Further, HAL has been used in a broad range of cognitive, neuropsychological, bilingual and language acquisition research (Bueno & Frenck-Mestre, 2002; Burgess & Lund, 1996, 1998; Li, Burgess, & Lund, 2000).

Modeling Concept Acquisition and Representation: The HAL Model

Developing a plausible methodology for representing the meaning of a word is central to any serious model of memory or language and has a long history in political science (Osgood, 1952). To this end, cognitive scientists have developed high-dimensional models of memory and language use. The models are high-dimensional in the sense that they represent word meanings in relation to other words used in a particular text or corpus of texts (Burgess & Lund, 2000; Landauer & Dumais, 1997; Lund & Burgess, 1996). The intuition behind these high-dimensional models is that the meaning of a word is learned from the contexts in which it is used, returning a number of meaning dimensions equal to the number of individual words used in a corpus, which can exceed 100,000. The benefit of the HAL model is that is able to learn the relationship of a word to other words as a function of the words' co-occurrence patterns providing an objective and transparent metric for measuring word similarity.

The concept acquisition process in HAL is a computationally simple model although the matrix of co-occurrence values is large-scale (e.g., its matrix is ultra-dimensional). The model's basic assumptions in computing meaning representation are:

1. The meaning of words develops through the usage of words in the context of other words;
2. The model learns word meanings from actual human language from a corpus of text;
3. A lexicon is constructed prior to building the matrix and these words serve as the row and column headers for the matrix;
4. The number of unique items in the lexicon determines the dimensionality of the model (e.g., 500 unique words = 500 context dimensions);
5. Humans associate words with other words that appear near a target word as the reader scans through text and meaning is encoded via word co-occurrences inside a window of working memory that moves along the text;
6. The window is relatively small;
7. The distances in the window are weighted relative to the proximity of a word to the words that precede it;
8. Co-occurrence values in each cell are normalized for word frequency in the overall corpus.

We use HAL to encode text using a window analogous to a person's working memory span during reading (see Gernsbacher, 1990). For the analyses here, we use a window width of 10 words: we assume that a person cognizes the word they are currently reading, along with the previous and later 10. In this sense, we allow these words to co-occur with a given target word in a window of weighted attention that moves through the text one word at a time, calculating co-occurrences at each step and storing the co-occurrence values in the memory matrix. In these analyses we only use the row values (co-occurrences that occur *prior* to the target). This rarely affects inferences in computational studies of language use and when it does, it is relevant to grammatical-syntactic constraints (Burgess & Lund, 1997a).

Another assumption is that as a particular word in the window is left further behind by the reader, its memory in relation to the current word fades. Consequently, we assign values inversely related to the distance between the two words in the window. For example, if two words were next to each other in a sentence they would be assigned a value of 10 (since a 10-word window is used), while a pair that is separated by nine words would only be assigned a value of 1. We also assume that if a word appears twice within a window of words, its relationship to the current word, the first word in the window, is best characterized by the sum of the distances. If a given word appears immediately before the current word (counted as a 10), as well as another word four words ahead of it (counted as a 6), the pair of words would be assigned a score of 16 in this particular window of words. A limited moving window and the weighting of co-occurrences within the window is one important component of the model that differentiates it from similar models that use a "bag-of-words" approach to encoding information such as LSA (Landauer & Dumais, 1997).

Measuring concepts using the HAL model

The original computational instantiation of the model was used to analyze a corpus of approximately 360 million words of diverse "conversations" in USENET, an Internet-based discussion system and precursor to forums like Google Groups or weblogs. Lund and Burgess (1996) built the USENET corpus to study the performance of the HAL model and the meaning of concepts in a text sample of English-language conversational speech. They pooled text from

USENET conversation forums and track lexical co-occurrence within a 10-word moving window. From the co-occurrences, a 140,000 dimensional context space was learned. This high-dimensional context or memory space is the word co-occurrence matrix. This high-dimensional semantic space is referred to as a "context" space since each vector element represents a symbol (usually a word) in the input stream of the text. Each symbol is part of the language context in the moving window.

Reviewing the basic methodology for computational studies using HAL, we begin with the development a matrix of word co-occurrence values for the lexical items in a given corpus. This matrix will then be divided into co-occurrence vectors for each word, which can be subjected to analysis for meaningful content. The product of this procedure is an *N*-by-*N* matrix, where *N* is the number of words in the vocabulary list.

This matrix contains significant amounts of information that can be used to simulate a variety of cognitive phenomena. The matrix is constrained by the vocabulary list (the lexicon) that is typically all words except very low frequency items that would not experience sufficient context to form a viable representation. A full co-occurrence vector for a word consists of both the row and the column for that word. These vectors can be viewed as the coordinates of points in a high-dimensional space, with each word occupying one point. Using this representation, we are able to compute the semantic similarity of two words as a function of differences between two words' co-occurrence vectors, using distance formulas from the Minkowski family:

$$d = \left(\sum_{i=1}^{n} |x_i - y_i| \right)^{1/p}$$

All analyses presented here use Euclidean distances (p = 2). The distance scores are scaled using Riverside Context Units, or RCUs, a metric scaled to the approximate range of word recognition times (in milliseconds), see Lund and Burgess (1996). Each element of a vector represents a coordinate in high-dimensional space for a word or concept, and distance metrics can be applied to these vectors that correspond to the similarity of the global contexts in which words are used for comparison across target words.

[Figure 1 about here]

In the memory matrix that was trained on the USENET corpus, the words *dog* and *cat* are closer to each other, i.e., they exhibit greater semantic similarity. Similarly, *street* and *road* are semantically similar to each other. However, *cat* and *street* have a longer semantic distance from each other. Similar concepts are computed as closer together than concepts that have less semantic (or contextual) similarity. The vectors that measure the meaning of words in relationship to other words can also be viewed graphically as can be seen in Figure 1. Sample words from the USENET corpus analysis (e.g., *dog*, *cat*) are shown with their accompanying 20 element vectors. Words representing similar concepts have similar vectors: A viewer can see the similarity of *dog* and *cat* as well as the similarity of *street* and *road* vectors. Each vector element has a continuous numeric value (the frequency normalized value from its matrix cell). A word's vector can be seen as a distributed representation; that is, each co-occurrence (vector element) is a coordinate in the high-dimensional semantic space (Rumelhart, Hinton, & McClelland, 1986).

Figure 2 illustrates examples of categorization effects that the HAL model has been used to investigate and uses a multidimensional scaling algorithm (MDS) which projects points from a high-dimensional space into a lower-dimensional space in a non-linear fashion. Figure 2a categorizes animals, foods, and geographic locations. Figure 2b illustrate a particular feature of HAL's meaning vectors, namely, that they can be used to model abstract concepts that have been problematic for representational theory (weather terms, proper names and emotional terms).

[Figure 2 about here]

Capturing the meaning of abstract concepts is crucial for investigating issues in political science and social constructs in political psychology. Modeling abstract concepts are important for a model that purports to be general in nature. Figure 2c illustrates the grammatical nature of word meaning (the categorization of nouns, prepositions, and verbs). The generalizability of the HAL model to capture grammatical meaning as well as more traditional semantic characteristics of words is an important feature of the model (Burgess, Livesay, & Lund, 1998; Burgess & Lund, 1997a) and was part of our motivation to refer to the high-dimensional space as a context space rather than a semantic space. These and other characteristics of word meaning that the model encodes has led us to rethink a number of assumptions about the dynamics of memory and concept acquisition (Burgess & Lund, 2000). The HAL model offers a clearly defined way to think about the nature of word relationship formation in the learning process and the relationship of basic associations to higher-order word meaning.

Demonstrations Using the HAL Model

The HAL model is versatile, and is applicable to a broad range of phenomena of interest to social science researchers. We show some of that range with two brief demonstrations of its potential: We explore the priming of racial stereotypes and the placement of politicians in ideological space. The purpose of these demonstrations is to show the HAL model's adaptability and that, using only text, it can reproduce patterns identified by researchers using drastically different means of observation. Both these demonstrations rely on a computational study of the 1992 USENET text corpus Lund and Burgess collected.

<u>Study 1: Priming racial stereotypes</u>

Semantic priming is a well-researched phenomenon among psychologists and many political psychologists. Priming reflects the idea that memory retrieval of a word will be facilitated if it is preceded by a related word (e.g., *cat* primes *dog*). The semantic distance between words computed using the HAL model is a much more accurate indicator of memory retrieval time (Burgess, 1998; Lund, Burgess, & Atchley, 1995; Lund, Burgess, & Audet, 1996), than are patterns of local co-occurrence. The conclusion from this is important. Global co-occurrence information better represents human learning and the semantic continuum than local co-occurrence or simple word frequency measures. Others have argued that HAL is a promising alternative to traditional models of semantics (Hutchison, 2003; Lucas, 2000). Merker (2004) argues that models like HAL and LSA are best suited for scaling up to large systems such as the brain and representing the neural substrate that supports memory for life history.

Another important issue with political concepts is their abstractness. Typically concepts in political science are going to be more abstract than stimuli used in basic cognition research. Is HAL capable of capturing the more subtle relationships that abstract words have? In a series of studies we showed that the distance metric directly corresponded to reaction time results in priming experiments for abstract and concrete words (Audet & Burgess, 1999), emotional words (Burgess & Lund, 1997b), and the difference between depressed and non-depressed individuals (Alison & Burgess, 2003). The semantic neighborhoods of these words can be seen as a *connotative* schema (e.g., *Beatles*: *original*, *band*, *song*, *album*, *first*, *British*). When subjects are presented with these neighborhoods they can reliably retrieve the explicit target word (20% of the subjects) or a very close semantic neighbor (Burgess et al., 1998) suggesting that the neighborhood contents has a cognitive parallel. The model is also capable of providing metrics that account for syntactic, thematic and grammatical effects unlike models that utilize a bag-of-words approach (Burgess, Livesay, & Lund, 1996; Burgess & Lund, 1997a) as well as problem-solving that hinges on semantic characteristics (Burgess & Lund, 2000).

There is considerable theoretical value to the notion of stereotypes. The distinction between automatic and controlled processing has shed considerable light on when a person may be aware of harboring a stereotype and to what extent they are able to keep that information private (Devine, 1989). As the population's awareness of racial stereotypes has grown and the cost of publicly holding such beliefs, the idea has emerged in social psychology that a person will try to hide their prejudices (Dovidio, Evans, & Tyler, 1986; Dovidio & Gaertner, 1977). Considerable recent research has investigated the extent to which these private beliefs can be detectable with semantic priming procedures and tasks that tend to dissociate the person from intent of the task (Burdein & Taber, 2004). These issues hold substantial significance for building a model of political semantics (Hutchings & Valentino, 2004; Sommers, Apfelbaum, Dukes, Toosi, & Wang, 2006; Valentino & Sears, 2005). Seemingly subtle racial information in political ads can be tactically used to be either prime or suppress racial attitudes (Valentino, Hutchings, & White, 2002). Policy oriented terms, such as "welfare" can racially bias information processing (Burdein & Taber, 2004).

A computational model of semantics would need to encode information pertaining to racial attitudes and behavior. An important contribution to political science from cognitive psychology has been theoretical and more recently computational models of meaning (McGraw, 2000; Taber & Timpone, 1996). In this experiment we assess the representations that are learned by the HAL model in a basic test of racial bias. The results will not directly bear on the issue of automatic and controlled processing since only representations of words are being used, however these representations could presumably provide the semantic structure for such processing if recent work in cognitive psychology is a guide (Burgess & Lund, 2000; Landauer & Dumais, 1997; Lund & Burgess, 1996; Lund et al., 1996; Rehder et al., 1998). The present goal is to use the stimuli derived by Greenwald (1998) that produced semantic priming independent of the person's stated beliefs (also see Banaji & Greenwald, 1994).

Method

Greenwald, et al. (1998) used a variety of word types in their experiments. In two experiments they capitalized on the semantic connotation of typical "black" and "white" names (e.g., *Adam*, *Betsy*, versus *Latisha*, *Tyrone*) in conjunction with words that have pleasant or unpleasant connotations (e.g., *love*, *friend*, versus *accident*, *vomit*). Previous research with the HAL model has shown that it is sensitive to proper name semantics (Burgess & Conley, 1998; Burgess & Conley, 1999; Burgess et al., 1996; Conley, Burgess, & Hage, 1999) and also ethnic characteristics of names (Willits et al., 2003). In this analysis, we used the words "*blacks*" and "*whites*" to prime the pleasant (positive) and unpleasant (negative) words they used.

In this analysis, we used slightly over half of positive and negative stimuli Greenwald et al. (1998) use (14 / 25 positive items; 16 / 27 negative items). Some items were not used because they were not in HAL's 70,000 item lexicon, were ambiguous (e.g., *diamond*), or were potentially policy related (e.g., *health*). If the analysis meets with success, the usage of "*blacks*" and "*whites*" as priming terms (rather than proper names) will result in a powerful tool for the testing of racial meaning issues with the model.

Results

The representation for the concept "*blacks*" is more similar to the meaning of the negative words (e.g., *sickness*, *grief*) than was the concept "*whites*" ($F(1, 16) = 334.89$, $p < .0001$). There was no reliable difference between "*whites*" and "*blacks*" with the set of positive words (e.g., *love*, *cheer*; $F(1, 13) = 0.126$, $p = 0.73$). There was also no reliable difference between "*blacks*" at the negative and positive levels of connotation ($F(1, 13) = 0.299$, $p = 0.59$)

or between “*blacks*” and “*whites*” at the negative and positive levels of connotation, respectively ($F(1, 13) = 1.66$, $p = 0.22$).

[Figure 3 about here]

Discussion

These results demonstrate a number of important effects for this area of modeling the semantics of social and political concepts. First, we see the utility of using the words “*blacks*” and “*whites*” to represent the concepts of black and white as they refer to ethnicity. Typically “*blacks*” and “*whites*” refer to ethnicity whereas “*white*” and “*black*” tend to be ambiguous. These two words, or from a high-dimensional modeling perspective, concepts occur in language contexts that support a host of negative connotations for blacks but not whites. A limitation to these results is that the language sample is USENET discussions where the composition of the writers is unknown. However, the HAL model has been used to model population differences by having appropriate language samples (specialized corpora) produced by a particular group of people; e.g., young, aged, Alzheimer’s patients (Conley & Burgess, 2000), or brain-damaged patients (Buchanan, Burgess, & Lund, 1996; Buchanan, Kiss, & Burgess, 2000). The distance metric (actually, the density of the semantic neighborhood – context density, in this case) predicts word recognition times for younger and older adults better than human word association norms (Buchanan, Westbury, & Burgess, 2001).

These results suggest that this approach may be exploited to explore racial stereotypes without having to involve human coding or tagging of test. Social and political psychology has been concerned with the implications of a person’s evaluative associations in memory as they pertain to racial and ethnic attitudes (Burdein, Lodge, & Taber, 2006; Burdein & Taber, 2004; Greenwald et al., 1998; Greenwald, Pickrell, & Farnham, 2002; Lodge, Taber, & Burdein, 2006).

High-dimensional memory modeling, treating text as language behavior and learning the connotative meaning of concepts is a step in the automatic determination of these associative networks.

Study 2: Ideological placement of U.S. Senators

Political scientists have developed a variety of ways to measure the ideological placement of political elites in policy space. Many of these rely on content analysis of text (e.g. Budge et al., 1987; Hill et al., 1997; Laver et al., 2003; Monroe & Maeda, 2004). Perhaps most prominent among ideal-point estimation techniques is the scaling of roll-call votes, pioneered by Poole and Rosenthal (1985). Poole and Rosenthal developed several ideological placement scores for members of the U.S. House and Senate – most voting behavior in Congress can be explained using their first dimension scores, which place these representatives on a left-right ideological continuum.

Methods

The HAL model used was based on the 320 million word 1992 USENET corpus. Senators whose names occurred at least 500 times and were not lexically ambiguous were selected. For example, we did not include U.S. Senator Arlen J. Specter (R-PA) in this analysis because of potential ambiguities associated with the word *specter*. We focus on eight names using this selection process. Full vectors (each with more than 68,000 weighted co-occurrence values) were extracted for the set of names and were submitted to a multi-dimensional scaling procedure to reduce the dimensionality in order to view the spatial relationships among the conceptual representations for the senators.

Results

The MDS for the set of senators can be seen in Figure 4. The Republicans (*Lugar*, *Gramm*, and *Dole*) form a conceptual space separate from the Democrats (*Dodd*, *Byrd*, *Kerry*, and *Kennedy*). Using Poole and Rosenthal's D-NOMINATE scores, we found one senator who could be considered liberal Republicans (*Hatfield*) and was in HAL's lexicon. In HAL's semantic space – a function of these USENET discussions – this senator is more similar to the Democrats we modeled.

In order to simulate the ideological dimension carried in the Poole and Rosenthal D-NOMINATE scores, distances in the HAL model were calculated for each senator's name and the words "*liberal*" and "*conservative*." This set of distances allowed the construction of a numeric scale that could be compared with the D-NOMINATE scores. Note that these distances were computed using the same vectors that contributed to the MDS solution in Figure 4. The HAL-based ideological placement scores for these U.S. Senators are strongly correlated with Poole and Rosenthal's D-NOMINATE scores ($r=.63$, $p<.1$). Both scales identify liberals with lower and negative numbers, and conservatives as higher, positive values.

Hill, et al. (1997) develop a measure of liberal-conservative ideological placement for U.S. Senators via content analysis of articles about the representatives' first campaigns for office from the *Chicago Tribune, Los Angeles Times, New York Times,* and the *Washington Post.* Thus, the authors measure ideology independent of behavior in the legislature. Their measure is highly correlated with other measures relying on roll call data, including Poole and Rosenthal's NOMINATE and scores from Americans for Democratic Action. The HAL-based ideological measure we develop is also correlated with the Hill, et al. content based measure ($r=.61$, $p=.1$).

Our intention is primarily to demonstrate that HAL reproduces elements of a semantic system meaningful in political terms. However, we also note that the measure we develop has two advantages over the Hill, et al. (1997) measure. First, the authors had to find text for newspaper articles from the first election campaign of the Senators. This prevented the authors from providing a score for some of the older members of the Senate (Hill et al., 1997, p. 1399). Perhaps more importantly, to perform their analysis, the authors had to train three human coders to examine the news stories and make judgments on the ideological attributions in the newspaper articles selected.

Any scholar who has attempted to code any text sample by hand can appreciate the time and expense in performing a large-scale content analysis. Our approach relies on a very large sample of text collected outside of the context of any particular election campaign. Consequently, we could envision creating a similarly valid measure using available samples of text from any point in American history. Further, the measurement technique we use requires no coder training and no human judgments with regard to content analysis itself. As a result of these advantages, we can provide a consistent measurement of ideology for any number of political elites quickly, and with a negligible increase in time expenditures.

Other Applications of the HAL Model to Political Science Research Questions

The HAL model has applicability beyond the demonstrations discussed here. We are particularly enthusiastic about the possibilities of pairing HAL-based computational studies with human-subject research to inform scholarship on political cognition, as well as potential for using the model to inform document categorization and the study of candidate ambiguity.

Political cognition

The HAL model's distance metric has a long history of accounting for a broad range of cognitive phenomena (Burgess et al., 1998; Burgess & Lund, 2000). The metric best corresponds to categorical semantic knowledge more so than simple associative relationships (Lund et al., 1995; Lund et al., 1996). The distances between positive and negative words and the words "*blacks*" and "*whites*" mimicked the type of social psychological bias that is found with humans (Banaji & Greenwald, 1994; Greenwald et al., 1998; Rudman, Greenwald, Mellott, & Schwartz, 1999). The model has a history of using the high-dimensional representations to account for abstract and emotional connotation (Audet & Burgess, 1999) and personality connotation (Hussain, Willits, & Burgess, 2006; Willits et al., 2003). The ability of the model to address racial issues is important. Racial bias is implicated in a wide set of issues in political science (Burdein & Taber, 2004).

The analysis of ideology and the Senators names demonstrates several very important characteristics of the conceptual representations that are learned by the HAL model. First, for the distance to correlate with the D-NOMINATE scores, the representations for the concepts of "*liberal*" and "*conservative*" simply must carry relevant ideological information. This information, in the form of weighted global co-occurrence values is acquired in the same way that representations are learned by the model in order to account for other cognitive phenomena (Buchanan et al., 2001; Burgess et al., 1998; Burgess & Lund, 1997a, 2000; Li et al., 2000). Ideology is an obvious problem for a computational semantic model that would be used for addressing issues in political science. However, this result is important for another very important reason. Developing semantic representations for proper names has been a long standing problem in the cognitive and information sciences. The HAL model has a long history of accounting for basic cognitive effects with humans and proper name retrieval (Burgess & Conley, 1998; Burgess & Conley, 1999; Burgess et al., 1996; Conley et al., 1999). The model has also accounted for the social aspects of proper names such as personality connotation (Willits et al., 2003). The ideological categorization of the senators and the correlation with the D-NOMINATE scores provides a substantial extension of the utility of the HAL model in accounting for proper name phenomena and addresses these important issues in the political context.

Categorization

HAL also has a great deal of promise as an complementary approach to some of the topic-identification and categorization problems identified by various political science researchers (Hopkins & King, 2007; Purpura, 2006; Purpura & Hilliard, 2006; Quinn, Monroe, Colaresi, Crespin, & Radev, 2006; Quinn et al., 2007; Shulman, Callan, Hovy, & Zavestoski, 2006) . While we have not yet applied HAL to the identification of topics in legislative actions or other governmental data, this is a promising avenue for further application of HAL. The model is particularly sensitive to how the meanings of concepts drift in this kind of data.

Computer scientists have used HAL to sort documents and have found the model performs better than Markov-based or relevance-based query models (Bruza & Song, 2002). It inspired a new algorithm for document clustering (Wang, Jia, & Yang, 2006) and improves document retrieval precision when implemented in Chinese (Cheong, Song, Bruza, & Wong, 2005). The HAL model has been used to investigate inductive scientific discovery with an implicit process (Bruza, Cole, Song, & Zeeniya, 2006; Bruza, Song, & McArthur, 2004) not unlike mediated priming (Livesay & Burgess, 1998). Bernhard (2006) investigated the automatic acquisition of semantic information using WordNet and the NCI Thesaurus and concluded that these more traditional approaches should be compared to models such as HAL and LSA. Representational or categorization systems have been influenced by statistical language research both in NLP (Manning & Schutze, 1999) and cognitive science (Murphy, 2002) – both discussing the HAL model in this context. HAL represents a useful alternative to traditional NLP semantic approaches such as Viegas (1999).

Ambiguity

While many scholars have developed approaches to placing parties and candidates in ideological and policy space, as discussed above, they have not studied strategic ambiguity. The modeling of ambiguity has been a challenge for high-dimensional models such as HAL and LSA until somewhat recently (Burgess, 2001; Landauer, 2001). The more ambiguous a word is the more diverse the history of the word is with other words. The word *Ray* could refer to a person or to a stream of light. What makes high-dimensional memory models excellent candidates for ambiguity research is that the learning process creates a vector of the word that is a history of the contexts in which it was used. This is critical because context sensitivity should enhance the sensitivity of the metrics to ambiguity detection.

With HAL, a new metric - context density - is the density of the co-occurrence values in a word's vector, and contributes nicely to determining the co-occurrence "features" of the word as well as their weight in representing the word's meaning and determining the relative dominance of the meaning which is influential in predicting the cognitive processing time of semantic ambiguity (Atchley, Burgess, Audet, & Arambel, 1996; Burgess, Tanenhaus, & Seidenberg, 1989; Burgess, Trueswell, Tanenhaus, & Garnsey, 1990; Simpson & Burgess, 1988) or syntactic ambiguity (Burgess & Hollbach, 1988; Burgess & Lund, 1994; Burgess, Tanenhaus, & Hoffman, 1994; Burgess et al., 1990). The syntactic effects occur with HAL because of the weighted distances in the co-occurrence window and are something that bag-of-words approaches will not obtain (Audet & Burgess, 1998). A characteristic of an ambiguous word is that by virtue of its multiple meanings it will have experienced more contexts than a less ambiguous word. Thus, the higher the proportion of non-zero cells in a word's co-occurrence vector, the more likely that word is to be ambiguous.

Conclusion

The HAL model, and other high-dimensional models of concept representation, complement other approaches to large-corpus text analysis and have a great deal of versatility in the investigation of political phenomena. We are able to systematically characterize the meaning of concepts and associations among objects in memory and apply these characterizations to a variety of political science research questions. Importantly, we are able to use HAL to reproduce findings from the study of racial stereotyping and the ideological placement of U.S. Senators. We show the usefulness of the HAL model as well as its practicality when compared to other methods. As we note, the model has additional applicability to the study of political cognition, campaign communication, and practical and administrative problems like document sorting.

The model not only enjoys this empirical face validity, but is also theoretically satisfying. It bases its measures on a theory of human cognition, concept acquisition, and representation. The model approximates the process of human learning from text. We emphasize the importance of context and experience in understanding social communication. Consequently, it provides a useful addition to the toolbox used to study text computationally which, at this early stage in its development, is dominated by techniques that take words out context as a simplifying assumption. We anticipate a profitable triangulation relying both on word frequency models and high-dimensional semantic approaches in the study of political communication.

References

Alison, J., & Burgess, C. (2003). Effects of chronic non-clinical depression on the use of positive and negative words in langague contexts. *Brain and Cognition, 53*, 125-128.

Atchley, R. A., Burgess, C., Audet, C., & Arambel, S. (1996). Timecourse, context effects, and the processing of lexical ambiguity in the cerebral hemispheres. *Brain and Cognition, 30*(3), 277-280.

Audet, C., & Burgess, C. (1998). Distinguishing between manner of motion and inherently directed motion verbs using a high-dimensional memory space and semantic judgments. Paper presented at the Cognitive Science Society. Madison, WI.

Audet, C., & Burgess, C. (1999). Using a high-dimensional memory model to evaluate the properties of abstract and concrete words. Paper presented at the Cognitive Science Society. Vancouver, BC.

Banaji, M., & Greenwald, A. G. (Eds.). (1994). *Implicit stereotyping and prejudice*. Hillsdale, NJ, England: Lawrence Erlbaum Associates, Inc.

Bernhard, D. (2006). Automatic acquisition of semantic relationships from morphological relatedness. *Lecture notes in computer science, 4139*, 121-132.

Bruza, P., Cole, R., Song, D., & Zeeniya, B. (2006). Towards operational abduction from a cognitive perspective. *Logic journal of the IGPL, 14*(2), 161-177.

Bruza, P., & Song, D. (2002). Inferring query models by computing information flow. Paper presented at the 11th International Conference on Information and Knowledge Management (CIKM 2002).

Bruza, P., Song, D., & McArthur, R. (2004). Abduction in semantic space: Towards a logic of discovery. *Logic journal of the IGPL, 12*(2), 97-109.

Buchanan, L., Burgess, C., & Lund, K. (1996). Overcrowding in semantic neighborhoods: Modeling deep dyslexia. *Brain and Cognition, 32*(2), 111-114.

Buchanan, L., Kiss, I., & Burgess, C. (2000). Phonological and semantic information in word and nonword reading in a deep dyslexic patient. *Brain and Cognition, 43*(1-3), 65-68.

Buchanan, L., Westbury, C., & Burgess, C. (2001). Characterizing semantic space: Neighborhood effects in word recognition. *Psychonomic Bulletin & Review, 8*(3), 531-544.

Budge, I., Robertson, D., & Hearl, D. (Eds.). (1987). *Ideology, strategy and party change: Spatial analysis of post-war election programmes in 19 democracies*. Cambridge, Great Britain: Cambridge University Press.

Bueno, S., & Frenck-Mestre, C. (2002). Rapid activation of the lexicon: A further investigation with behavioral and computational results. *Brain & Language. Special Issue: Mental lexicon II, 81*(1-3), 120-130.

Burdein, I., Lodge, M., & Taber, C. (2006). Experiments on the automaticity of political beliefs and attitudes. *Political Psychology, 27*(3), 359-371.

Burdein, I., & Taber, C. (2004). What are they really thinking? Implicit measures of ideology and racism. Paper presented at the The Midwest Political Science Association. Chicago, IL.

Burgess, C. (1998). From simple associations to the building blocks of language: Modeling meaning in memory with the hal model. *Behavior Research Methods Instruments & Computers, 30*(2), 188-198.

Burgess, C. (2001). Representing and resolving semantic ambiguity: A contribution from high-dimensional memory modeling. In D. S. Gorfein (Ed.), *On the consequences of meaning selection: Perspectives on resolving lexical ambiguity.* (pp. 233 - 290). Washington, DC: American Psychological Association.

Burgess, C., & Conley, P. (1998). Developing semantic representations for proper names. Paper presented at the Cognitive Science Society, Madison, WI.

Burgess, C., & Conley, P. (1999). Representing proper names and objects in a common semantic space: A computational model. *Brain and Cognition, 40*(1), 67-70.

Burgess, C., & Hollbach, S. C. (1988). A computational model of syntactic ambiguity as a lexical process. Paper presented at the Cognitive Science Society. Montreal, CA.

Burgess, C., Livesay, K., & Lund, K. (1996). Modeling parsing constraints in high-dimensional semantic space: On the use of proper names. Paper presented at the Cognitive Science Society. University of California, San Diego.

Burgess, C., Livesay, K., & Lund, K. (1998). Explorations in context space: Words, sentences, discourse. *Discourse Processes, 25*(2-3), 211-257.

Burgess, C., & Lund, K. (1994). Multiple constraints in syntactic ambiguity resolution: A connectionist account of psycholinguistic data. In *Proceedings of the cognitive science society* (pp. 90-95). Hillsdale, N.J.: Lawrence Erlbaum Associates.

Burgess, C., & Lund, K. (1996). Modeling the cerebral asymmetries of semantic memory in high-dimensional space: HAL's brain. *Brain and Cognition, 30*(3), 6-6.

Burgess, C., & Lund, K. (1997a). Modeling parsing constraints with high-dimensional context space. *Language and Cognitive Processes, 12*(2-3), 177-210.

Burgess, C., & Lund, K. (1997b). Representing abstract words and emotional connotation in a high-dimensional memory space. Paper presented at the Cognitive Science Society. Palo Alto, CA.

Burgess, C., & Lund, K. (1998). Modeling cerebral asymmetries in high-dimensional semantic spaces. In M. Beeman & C. Chiarello (Eds.), *Right hemisphere language comprehension: Perspectives from cognitive neuroscience* (pp. 215-244). Hillsdale, NJ: Lawrence Erlbaum Associates.

Burgess, C., & Lund, K. (2000). The dynamics of meaning in memory. In E. Deitrich & A. B. Markman (Eds.), *Cognitive dynamics* (pp. 117-156). Mahwah, NJ: Lawrence Erlbaum Associates, Publishers.

Burgess, C., Tanenhaus, M. K., & Hoffman, M. (1994). Parafoveal and semantic effects on syntactic ambiguity resolution. Paper presented at the Cognitive Science Society, Atlanta, GA.

Burgess, C., Tanenhaus, M. K., & Seidenberg, M. S. (1989). Context and lexical access - implications of nonword interference for lexical ambiguity resolution. *Journal of Experimental Psychology-Learning Memory and Cognition, 15*(4), 620-632.

Burgess, C., Trueswell, J., Tanenhaus, M. K., & Garnsey, S. (1990). Interaction of semantic, syntactic, and visual factors in syntactic ambiguity resolution. *Bulletin of the Psychonomic Society, 28*(6), 522-522.

Cheong, P., Song, D., Bruza, P., & Wong, K.-F. (2005). Information flow analysis with chinese text. *Lecture notes in computer science, 3248*, 100-109.

Conley, P., & Burgess, C. (2000). A computational approach to modeling population differences. *Behavior Research Methods Instruments & Computers, 32*(2), 274-279.

Conley, P., Burgess, C., & Hage, D. (1999). Large-scale databases of proper names. *Behavior Research Methods Instruments & Computers, 31*(2), 215-219.

Devine, P. G. (1989). Stereotypes and prejudice: Their automatic and controlled components. *Journal of Personality and Social Psychology, 56*(1), 5-18.

Dovidio, J. F., Evans, N., & Tyler, R. B. (1986). Racial stereotypes: The contents of their cognitive representations. *Journal of Experimental Social Psychology, 22*(1), 22-37.

Dovidio, J. F., & Gaertner, S. L. (1977). The subtlety of white racism, arousal and helping behavior. *Journal of Personality and Social Psychology, 35*, 691-707.

Gernsbacher, M. A. (1990). *Language comprehension as structure building*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. K. (1998). Measuring individual differences in implicit cognition: The implicit association test. *Journal of Personality and Social Psychology, 74*(6), 1464-1480.

Greenwald, A. G., Pickrell, J. E., & Farnham, S. D. (2002). Implicit partisanship: Taking sides for no reason. *Journal of Personality and Social Psychology, 83*(2), 367-379.

Hart, R. P. (2000). *Campaign talk: Why elections are good for us*. Princeton, NJ: Princeton University Press.

Hill, K. Q., Hanna, S., & Shafqat, S. (1997). The liberal-conservative ideology of U.S. Senators: A new measure. *American Journal of Political Science, 41*(4), 1395-1413.

Hopkins, D., & King, G. (2007). Extracting systematic social science meaning from text. Unpublished Typescript. Harvard University.

Hussain, M., Willits, J., & Burgess, C. (2006). Proper name connotations: Testing explicit and implicit use of stereotypes. Paper presented at the Psychonomic Society. Houston, TX.

Hutchings, V. L., & Valentino, N. A. (2004). The centrality of race in American politics. *Annual Review of Political Science, 7*, 383-408.

Hutchison, K. A. (2003). Is semantic priming due to association strength or feature overlap? A microanalytic review. *Psychonomic Bulletin & Review, 10*(4), 785-813.

Landauer, T. K. (2001). Single representations of multiple meanings in latent semantic analysis. In D. S. Gorfein (Ed.), *On the consequences of meaning selection: Perspectives on* *resolving lexical ambiguity.* (pp. 217 - 232). Washington, DC: American Psychological Association.

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*, 211-240.

Laver, M., Benoit, K., & Garry, J. (2003). Extracting policy positions from political texts using words as data. *American Political Science Review, 97*(2), 311-331.

Li, P., Burgess, C., & Lund, K. (2000). The acquisition of word meaning through global co-occurrence. Paper presented at the 30th Child Language Research Forum, Stanford, CA.

Livesay, K., & Burgess, C. (1998). Mediated priming in high-dimensional semantic space: No effect of direct semantic relationships or co-occurrence. *Brain and Cognition, 37*(1), 102-105.

Lodge, M., Taber, C. S., & Burdein, I. (2006). Experiments on the automaticity of political beliefs and attitudes. *Political Psychology, 27*(3), 359-372.

Lucas, M. (2000). Semantic priming without association: A meta-analytic review. *Psychonomic Bulletin & Review, 7*(4), 618-630.

Lund, K., & Burgess, C. (1996). Producing high-dimensional semantic spaces from lexical co-occurrence. *Behavior Research Methods Instruments & Computers, 28*(2), 203-208.

Lund, K., Burgess, C., & Atchley, R. A. (1995). Semantic and associative priming in high-dimensional semantic space. Paper presented at the Cognitive Science Society, Pittsburgh, PA.

Lund, K., Burgess, C., & Audet, C. (1996). Dissociating semantic and associative word relationships using high-dimensional semantic space. Paper presented at the Cognitive Science Society, University of California, San Diego.

Manning, C. D., & Schutze, H. (1999). *Foundations of statistical natural language processing*. Cambridge, MA.: The MIT Press.

McGraw, K. M. (2000). Contributions of the cognitive approach to political psychology. *Political Psychology, 21*(4), 805-832.

Merker, B. (2004). Cortex, countercurrent context, and dimensional integration of lifetime memory. *Cortex, 40*(3), 559-576.

Monroe, B. L., & Maeda, K. (2004). Talk's cheap: Text-based estimation of rhetorical ideal-points. Paper presented at the 21st Annual Summer Meeting, Society for Political Methodology. Stanford University.

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

Osgood, C. E. (1952). The nature and measurement of meaning. *Psychological Bulletin, 49*, 197-237.

Pennebaker, J., & Graybeal, A. (2001). Patterns of natural language use: Disclosure, personality, and social integration. *Current Directions in Psychological Science, 10*, 90-93.

Poole, K. T., & Rosenthal, H. (1985). Recent developments in analytical models of voting in U.S. Congress. *American Journal of Political Science, 29*, 357-384.

Purpura, S. (2006, July 20-23, 2006). Sensitivity analysis in automated content analysis using a political party platform example. Paper presented at the Society for Political Methodology, University of California, Davis.

Purpura, S., & Hilliard, D. (2006). Automated classification of congressional legislation. Paper presented at the Seventh International Conference on Digital Government Research, San Diego, CA.

Quinn, K. M., Monroe, B. L., Colaresi, M., Crespin, M. H., & Radev, D. R. (2006). An automated method of topic-coding legislative speech over time with application to the 105th-108th U.S. Senate. Paper presented at the Summer Political Methodology annual meeting. Davis, California

Quinn, K. M., Monroe, B. L., Colaresi, M., Crespin, M. H., & Radev, D. R. (2007). How to analyze political attention with minimal assumptions and code. Unpublished Typescript. Harvard University.

Rehder, B., Schreiner, M. E., Wolfe, M. B. W., Laham, D., Landauer, T. K., & Kintsch, W. (1998). Using latent semantic analysis to assess knowledge: Some technical considerations. In P. W. Foltz (Ed.), *Discourse processes: A multidisciplinary journal* (Vol. 25, pp. 337-354). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Rudman, L. A., Greenwald, A. G., Mellott, D. S., & Schwartz, J. L. K. (1999). Measuring the automatic components of prejudice: Flexibility and generality of the implicit association test. *Social Cognition, 17*(4), 437-465.

Rumelhart, D. E., Hinton, G. E., & McClelland, J. L. (1986). A general framework for parallel distributed processing. In D. E. Rumelhart & J. L. McClelland (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition* (Vol. 1, pp. 45-76). Cambridge, MA: MIT Press.

Schonhardt-Bailey, C. (2005). Measuring ideas more effectively: An analysis of bush and kerry's national security speeches. *PS-Political Science & Politics, 38*(4), 701-711.

Shulman, S., Callan, J., Hovy, E., & Zavestoski, S. (2006). Language processing technology for electronic rulemaking: A project highlight. Paper presented at the Seventh National Conference on Digital Government Research.

Simon, A. F., & Xenos, M. (2004). Dimensional reduction of word-frequency data as a substitute for intersubjective content analysis. *Political Analysis, 12*(1), 63-75.

Simpson, G. B., & Burgess, C. (1988). Implications of lexical ambiguity resolution for word recognition and comprehension. In S. L. Small, G. W. Cottrell & M. K. Tanenhaus (Eds.), *Lexical ambiguity resolution: Perspectives from psycholinguistics, neuropsychology, and artificial intelligence* (pp. 311-329). San Mateo, CA: Morgan Kaufmann Publishers.

Slapin, J. B., & Proksch, S.-O. (2007). A scaling model for estimating time-series party positions from texts. Unpublished Typescript. University of California, Los Angeles.

Sommers, S. R., Apfelbaum, E. P., Dukes, K. N., Toosi, N., & Wang, E. J. (2006). Race and media coverage of hurricane katrina: Analysis, implications, and future research questions. *Analyses of Social Issues and Public Policy (ASAP), 6*(1), 39-55.

Starr, K. W. (1998). *Referral from Independent Counsel Kenneth W. Starr in conformity with the requirements of Title 28, United States code, Section 595(c).* (No. 105-310): U.S. Government Printing Office.

Stone, P. J., Dunphy, D. C., Smith, M. S., & Ogilvie, D. M. (1966). *The general inquirer: A computer approach to content analysis*. Cambridge, MA: The M.I.T. Press.

Taber, C. S., & Timpone, R. J. (1996). *Computational modeling*. Newbury Park, CA: Sage.

Valentino, N. A., Hutchings, V. L., & White, I. K. (2002). Cues that matter: How political ads prime racial attitudes during campaigns. *American Political Science Review, 96*(1), 75-90.

Valentino, N. A., & Sears, D. O. (2005). Old times there are not forgotten: Race and partisan realignment in the contemporary south. *American Journal of Political Science, 49*(3), 672-688.

Viegas, E. (Ed.). (1999). *Breadth and depth of semantic lexicons* (Vol. 10). Dordrecht: Kluwer Academic Publishers.

Wang, Y. H., Jia, Y., & Yang, S. (2006). Short documents clustering in very large text databases. *Lecture notes in computer science, 4256*, 83-93.

Willits, J., Burgess, C., Andrews, C., Khan, A., Hatsopoulos, A., & Conley, P. (2003). Measuring personality connotation in proper names using high-dimensional memory models. *Psychonomic Society Abstracts, 8*, 95.

Figure 1. Sample 20-element word vectors for four words computed using HAL

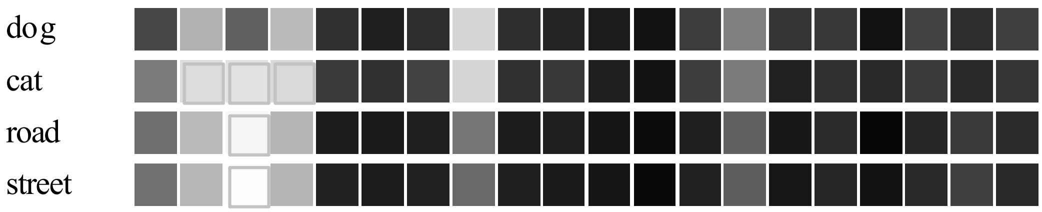

dog	7.56	44.90	13.42	49.46	3.75	1.59	3.15	64.86	3.26	2.08
cat	18.04	55.64	58.95	54.21	4.36	2.95	5.26	52.27	3.05	3.96
road	21.98	61.65	108.06	58.48	1.85	1.57	2.03	24.71	1.76	1.95
street	24.46	64.30	124.61	63.81	2.49	1.67	2.25	21.88	2.15	1.62

Figure 1. Sample 20-element word vectors for four words. Each vector element has a continuous value (the normalized value from its matrix cell) and is gray-scaled to represent the normalized value with black corresponding to zero. Below the gray-scaled vectors are the normalized numeric representations for the first 10 vector elements.

Figure 2. Two-dimensional multidimensional scaling solutions for three types of words.

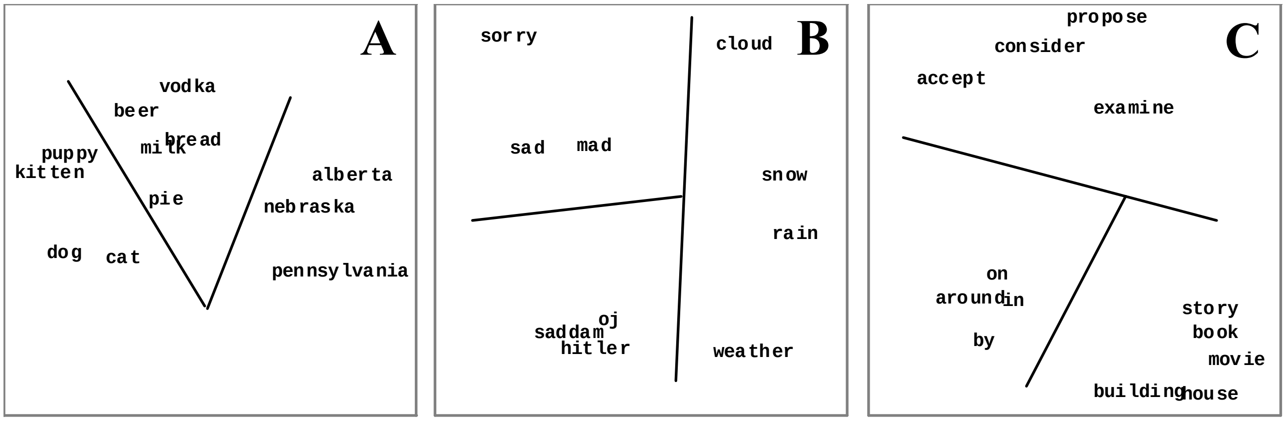

Figure 2. Two-dimensional multidimensional scaling solutions for : (A) common nouns, (B) abstract words, and (C) grammatical categories.

Figure 3. The interaction of race and emotional word connotation

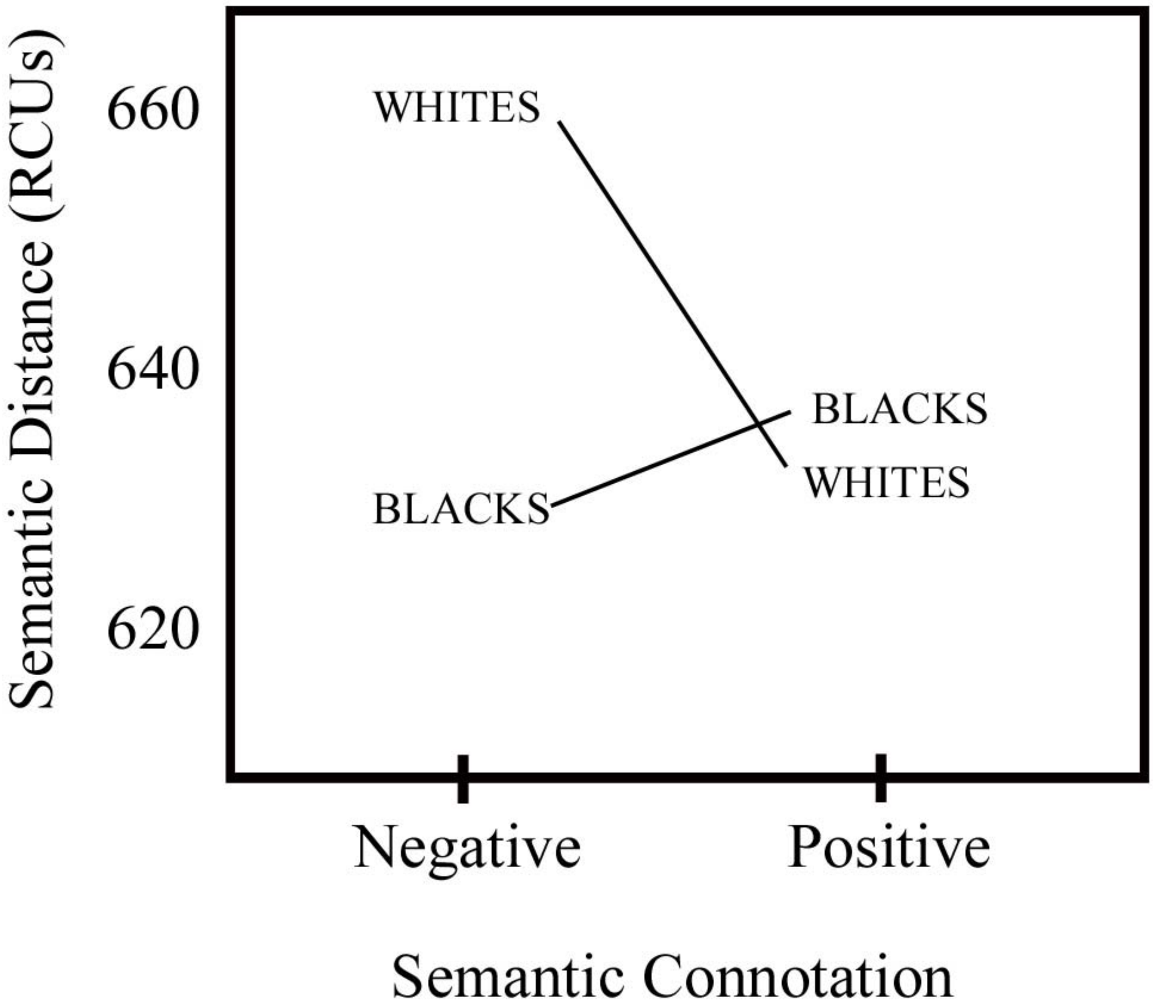

The interaction of race and emotional word connotation using the distance metric from the HAL model.

Figure 4. Distinguishing U.S. Senators by ideological orientations, using USENET text.

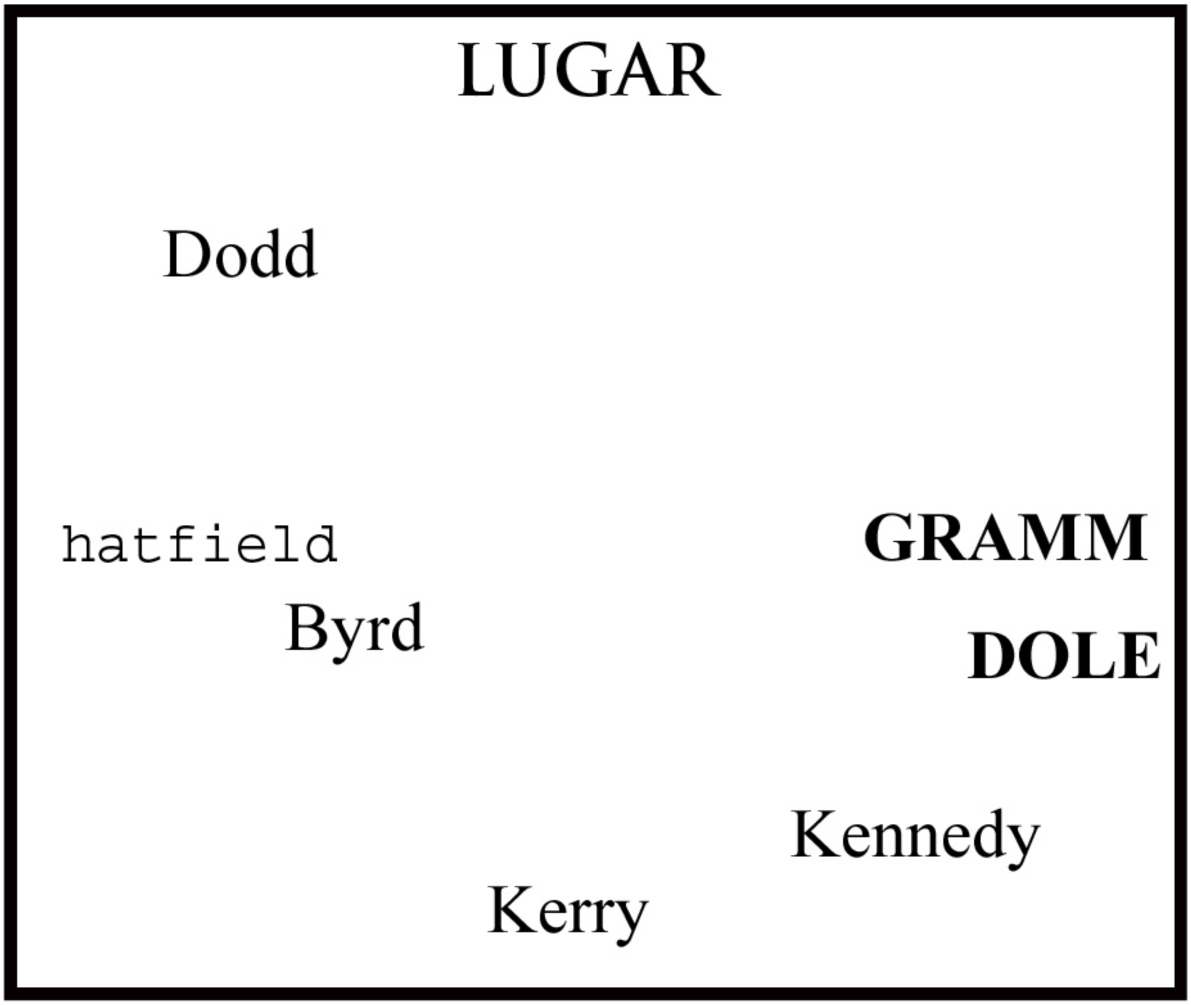

Two-dimensional multi-dimensional scaling solution for three ideological categories of U.S. Senator names. USENET contributor discussion allow us to discriminate between Republicans (Lugar, Gramm, Dole), Democrats (Dodd, Byrd, Kerry, Kennedy) and the liberal Republican (Hatfield) using the HAL model.