

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Developing Semantic Representations for Proper Names

Permalink

<https://escholarship.org/uc/item/8c00p5j0>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 20(0)

Authors

Burgess, Curt

Conley, Patrick

Publication Date

1998

Peer reviewed

Developing Semantic Representations for Proper Names

Curt Burgess (curt@doumi.ucr.edu)

Department of Psychology
University of California, Riverside
Riverside, CA 92521

Patrick Conley (patrick@cassandra.ucr.edu)

Department of Psychology
University of California, Riverside
Riverside, CA 92521

Abstract

A series of simulations using the HAL model of memory demonstrates that word representations from the model can be used to categorize a variety of common and famous proper nouns, cities, and states. The internal semantics of famous proper names provides a richer set of meaning constraints than do the neighborhoods of common proper names. Retrieval errors with names may be due to this difference in the neighbors and the density of these neighborhoods. A very salient constraint in common proper name semantics is gender.

Proper names are the Rodney Dangerfield of linguistics. They don't get no respect. -- McDonald (1996)

The investigation of proper name semantics has been a peripheral domain in all areas of cognitive science. The domains of cognitive neuropsychology and aging account for most proper name research (see Cohen & Burke, 1993). The paucity of work on proper names is notable when one considers that proper names are difficult to learn and remember and would, on the surface, seem to be an inviting area of investigation. The importance of proper names becomes foregrounded when one considers how often, during the course of a day, a person is required to retrieve a proper name when interacting with another person, and the necessity of retrieving other information associated with a proper name.

Memory Phenomena of Proper Names

Proper names are difficult to learn compared to other personal information. When subjects learned different kinds of biographical information about fictitious persons, Cohen and Faulkner (1986) found that proper name retrieval was more problematic compared to memory for other personal data such as hobbies or occupations. Retrieval failure with proper names is likely due to the lack of specific semantic information in the name's representation. Cohen and Burke (1993) argue that, without this semantic depth, there are few cues to use in the retrieval process. For example, if a person tries to retrieve *chair* and instead retrieves *seat*, the result would likely still be useable in the retrieval context. Although an analogous error can occur with proper names,

the outcome is quite different. Retrieving *Dave* when *Dan* was intended results in an error with more serious consequences. Naming problems are also frequent in aphasia. Semenza (1995) discusses cases in which individuals have demonstrated anomia for proper names as well as a case in which proper name information was spared.

Models of Proper Names

Cases of dual dissociation can be used to argue for a dissociation between memory for proper names and memory for semantic information in an unimpaired system such as in the models of Burton and Bruce (1993). Most models of proper names share the assumption of a separate memory store for proper names. Burton and Bruce offer a new approach with an interactive activation model and argue that names are difficult to retrieve not because they are "meaningless," but because they are unique -- there is just one Bill Clinton.

There has been very little discussion of the representational nature of proper names (Cohen & Burke, 1993). The purpose of this paper is to present a series of simulations and experiments that detail what the Hyperspace Analogue to Language (HAL) model of memory can offer to representational issues with proper names. We will present four experiments -- the first two will show how HAL's semantics can be used for the categorization of proper names and other concepts; then we turn to a description of the internal semantics of common and famous proper names; and finally a discussion of gender in proper names.

The HAL Model of Memory

HAL is a model that develops word meanings from a ~300 million word corpus of Usenet text by utilizing a global co-occurrence learning algorithm. HAL stores weighted co-occurrence information for the 70,000 most frequent symbols. This set of weighted co-occurrences forms a matrix of co-occurrences for words in the context of other words. This matrix results in a set of word vectors that serve as coordinates for the word in a 140,000 high-dimensional meaning space. The contextual similarity of the words is inversely proportional to the distance between the words in the hyperspace (Burgess & Lund, 1997; Lund & Burgess, 1996).

The HAL model has been used to investigate a broad range of cognitive phenomena, including semantic priming effects and the dissociation between semantic and associative priming (Lund, et al., 1996). Although most work with HAL has been in the domain of traditional semantic and categorization research, the word vector representations encode meaning at a general level, not just what is typically considered semantic. The strongest evidence for this is how HAL's representations encode grammatical categorical information (Burgess & Lund, 1997). As a result, HAL's vector representations are best characterized as context vectors, rather than semantic vectors, since they encode the contexts in which words appear and make use of this information in a variety of ways (Landauer & Dumais, 1997).

Experiment 1: Proper Names as a Meaning Class

The goal of Experiment 1 is to show that HAL's basic ability for categorization extends to proper names. In this simulation, different groups of words are subjected to multidimensional scaling (MDS) in order to show that the interword distances in the high-dimensional meaning space can provide a basis for categorization.

Methods

Each vector is a set of coordinates in high-dimensional Euclidean space. Co-occurrence vectors were extracted for all words used in this experiment; the MDS solution is shown in Figure 1a.

Results

This MDS projects points in a 140,000 high-dimensional space into a two-dimensional space in a non-linear fashion that attempts to preserve the distances between points. Visual inspection suggests that the three categories of words were separated in the meaning space. Since the dimensionality reduction is substantial, an analysis of variance (ANOVA) of the distances in the high-dimensional space is conducted to provide another estimate of categorization. Lines have been drawn in the figures to make even more apparent the separation of these different classes of items; the ANOVAs are important to substantiate this visual (and theoretical) separation. The ANOVA compares the intragroup distances to intergroup distances. In this way, one can determine in the high-dimensional space if one group of words is differentiated from the others by the vector representations. The proper names are separate from the other two groups of items, $F(1, 42) = 109.71, p < .0001$. Likewise, verbs are in their own space, $F(1, 42) = 50.65, p < .0001$; as are the nouns, $F(1, 42) = 143.49, p < .0001$. One can observe internal semantics of the items; similar verbs are close together, although *liked* and *loved* are distinct from *shoved* and *pushed*, $F(1, 4) = 22.08, p < .009$. There is a suggestion that male proper names occupy a distinct space from the female names and this gender issue will be

explored in more detail in Experiment 4.

Discussion

The results from this simulation support the contention that words with similar meanings share similar areas in the high-dimensional space. This conclusion holds for what would be considered more semantic relationships as well as grammatical categories; similar to earlier results (Burgess & Lund, 1997). The novel finding is that proper names seem to occupy a space distinct from other nouns. We have not yet shown that proper names have any meaningful internal semantics. In the next experiment we investigate the extent that HAL's vector representations can make further distinctions among proper names and in Experiment 3 we explore the question about the internal semantics and the relationship between a name's semantic characteristics and the cognitive effects for proper names reported in the literature.

Experiment 2: Categorization of Diverse Types of Proper Names

One domain in which proper names have been extensively investigated is in text-based information retrieval (IR). However, proper names can extend to a very diverse group of concepts: cities, governments, states, famous persons, diseases, religions, and so on. Approaches to this proper name categorization have been varied but tend to include what McDonald (1996) refers to as internal and external evidence. Internal evidence involves the form of a name, which is usually recorded in some database (e.g., *Cornhuskers* stored in a database of great football teams). External evidence uses the context surrounding a name (e.g., *President Clinton* identifies this particular Clinton as the President). Compilations of databases such as these can provide the front end to a sophisticated tool that can, in conjunction with part-of-speech taggers, parsers, morphological analyzers, text structurers, and other components, generate a very reasonable categorization of a diverse set of proper names. One of the more sophisticated of these can correctly categorize 96% of the names it encounters (Paik, Liddy, Yu, & McKenna, 1996). Its unclear, though, if these approaches will correspond well to the cognitive processing that takes place during name retrieval. In this experiment we see if a more diverse set of proper names can be categorized. The items used in this experiment were motivated by the range of proper nouns described by Paik et al.

Methods and Results

Co-occurrence vectors were extracted for four classes of proper names: cities, states, famous proper names, and common proper names. A MDS solution was computed (see Figure 1b). Visual inspection suggests that these diverse categories of proper names were differentiated. Cities were separated in this procedure from the other groups, $F(1, 70) = 49.13, p < .0001$; as were the states, $F(1, 70) = 28.8, p <$




Figure 1 : (a) Names, nouns, and verbs, (b) Types of proper nouns

.0001. Although cities and states share many features and are found in similar contexts, cities were separate from the states, $F(1, 26) = 46.87, p < .0001$. Famous proper names reside in their own high-dimensional space, $F(1, 106) = 32.89, p < .0001$; as do the common proper names, $F(1, 88) = 65.49, p < .0001$. Furthermore, the famous names are separate from the common names, $F(1, 58) = 99.79, p < .0001$. Further internal semantics is apparent as well. The co-occurrence vectors for the two presidents place the words very close. Likewise, murderers (or those accused of murder) share a similar space.

Discussion

Once again, the high-dimensional meaning space appears capable of providing the basic contextual information for the categorization of proper names. Although there is similarity along some dimensions for certain items, for example, *Reagan* and *Hitler* were both politicians; they are different within that domain. The results are encouraging for further consideration in that these vector representations might be a useful adjunct for information retrieval engines as they seem able to approach some of the same domains as more traditional approaches with considerably less top-down overhead.

The within-group similarity of some of these items (e.g., Reagan and Clinton) suggests that the internal semantics of these words may provide useful information related to cognitive issues in proper name retrieval. It is this matter that the next experiment addresses.

Experiment 3: Semantic Neighborhood Characteristics of Proper Names

Experiments 1 and 2 demonstrated that HAL's word representations provide a considerable information that can contribute to categorization. This "information" is a set of contextual cues for a word, i.e., all the weighted co-occurrence information that a word has associated with it as the moving window slides along the text. It is not difficult to imagine that this contextual information might allow for the kind of categorization presented in these experiments;

but, at the same time, perhaps not really provide further information that could provide support or insights into the cognitive processes involved in proper name retrieval and representation. In Experiment 3, we introduce the notion of the semantic neighborhood and evaluate its potential role in the cognitive theory of name representation. A word's vector elements are other symbols in the original text. The value for a particular element corresponds to the weighted co-occurrence for that word in the weighted context of the other symbol (usually a word). Each of these vector elements is a dimension in the high-dimensional space. Any word in HAL's 70,000 item vocabulary can be isolated in the high-dimensional context space. In that space, a word will have other words that vary in distance from it. Words that are close are referred to as "neighbors." In previous research, Burgess, Livesay, and Lund (1998) have shown that these neighborhoods provide sufficient cues to allow human subjects to distinguish between similar words.

To illustrate the notion of semantic neighborhoods, the closest 12 neighbors to two words are listed in Table 1. The neighborhoods differ in a number of respects. *Book's* neighbors include words that are very close in meaning to book (*story, books, paper, issue*). Other neighbors are similar in that they are other forms of informative media (*movie, film*), and others are more peripherally related (*game, new, recent, original, word, series*). The nature of the neighborhood for *John* is different than *Book's*. *John's* neighbors are simply other common proper names, and, in fact, are all male names. Famous proper names have neighborhoods that have a richer set of cues, more like objects (see Table 2).

There are two characteristics of these neighborhoods that suggest that HAL's neighborhoods may provide cognitively plausible information that would facilitate the representational modeling of proper names. Consider what happens in the case of a naming error. Presumably, the error occurs because a cognitive process (spreading activation, a semantic cue) was imprecise to some degree. If that happens for *book*, and another nearby word is retrieved, it is likely that the error will be either something very close in meaning to book, or some cue that might provide for the generation of other relevant cues. Conversely, if there is some error in a proper name, the wrong name will be retrieved. Proper name neighborhoods lack the semantic depth to provide any rich cues for retrieval if required. This corresponds nicely to most accounts of proper name retrieval errors in the literature (Cohen & Burke, 1993). The second characteristic that could help explain these types of errors has to do with the density of the neighborhoods. The numbers listed beside each neighbor in Table 1 are distances (in RCUs; see Lund & Burgess, 1996, for details) between the word and the neighbor. Note that the distance at the end of these particular neighborhoods differ by approximately 50 RCUs; the proper name neighborhood (at least in this case) is denser than the neighborhood for the common object, *book*. The implication of this is clear; any error in the retrieval process

is more likely to result in retrieval of a neighbor. The nearest neighbor of *John* is *David*, only 10 RCUs away; whereas the nearest neighbor of *book* is *story*, 17 RCUs away. In this experiment, we test to see if this is a general result for both high and low frequency matched nouns and proper names similar to those shown in Table 1.

Table 1: Semantic Neighbors of *Book* and *John*

Neighbors of <i>book</i>		Neighbors of <i>john</i>	
story	= 343	david	= 299
game	= 360	peter	= 309
movie	= 368	paul	= 311
new	= 368	mike	= 324
film	= 382	steve	= 334
books	= 384	mark	= 334
paper	= 394	james	= 335
recent	= 395	jim	= 337
original	= 396	robert	= 342
issue	= 398	tom	= 344
series	= 398	richard	= 347
word	= 399	michael	= 349

Table 2: Neighborhoods for Politician and Rock Group

Neighbors of: Reagan		Neighbors of: Beatles	
clinton	poor	original	british
republicans	congress	band	best
bush	american	movie	x-files
democrats	people	song	songs
republican	japanese	first	last
bill	president	main	album
americans		lyrics	

Methods

20 high and 20 low frequency proper names were frequency matched with 40 nouns. All words were taken from the Usenet corpus, and their frequency was a count of times the words appeared in the HAL corpus. This frequency matching was performed to control for the finding that higher word frequencies result in denser word neighborhoods.

Results

Both high and low frequency names had higher neighborhood densities than their frequency-matched noun counterparts, $t(19) = 8.12, p < .001$; $t(19) = 3.12, p < .01$. This result demonstrates that even when names and nouns appear the same number of times in the corpus, proper names have denser word neighborhoods in high-dimensional space.

Discussion

The analysis of the proper name and common object neighborhoods provides several insights into the difficulty commonly seen with proper names. Neighbor similarity with common objects reduces the potential impact of an error. An error with a proper name will almost certainly

result in the wrong name. This is exacerbated by the difference in the densities of these two types of neighborhoods. In Experiments 1 and 2 there was some indication in an effect of gender -- male and female names appear to occupy separate spaces. Furthermore, the neighborhoods of male and female names differ in their composition of male and female names. In Experiment 4 we explore the nature of this possible gender effect.

Experiment 4a: Gender Categorization in High Dimensional Space

The previous results provide evidence that proper names occupy their own locale in high-dimensional space. But what of categorizations within the pure name space itself? Specifically, we were interested in determining what effect the gender of the proper names has on that name's placement within high-dimensional space.

Methods

A set of high-frequency male and female names was drawn from the corpus, and co-occurrence vectors were computed for the names. An MDS was performed on the co-occurrence vectors to determine whether names would categorize in high-dimensional space based on gender.

Results

Analysis of the MDS (Figure 2a) demonstrates that the model distinguished well between male and female names, clustering each in its own location in high dimensional space. In addition to the MDS solutions demonstrating gender-based discriminability, an ANOVA was performed on the intragroup and intergroup name distances to provide inferential support for the MDS model. In this analysis, male names were differentiated from female names, $F(1,188) = 261.80, p < .001$. The co-occurrence vectors provide evidence that these male and female proper names are clustering in distinguishable groups in high-dimensional space. However, there is a significant difference in frequency between male and female names in the corpus. Since the proper names used were rank-ordered and the highest frequency male and female names used, this difference in frequency was not taken into effect. Since frequency can play a strong role in co-occurrence vectors, it is important to determine whether the gender-based proper name clustering in high dimensional space is merely a frequency effect or due to gender discriminability in the proper name co-occurrence vectors.

Experiment 4b: Frequency-Matched Gender Categorization

The previous simulation supplied evidence that names were categorized by gender in high-dimensional space. However, whether this was due to the greater frequency of the male proper names or was the result of gender information contained within the proper name co-occurrence vectors was




Figure 2: (a) Proper name gender categorization, (b) Frequency controlled proper name gender categorization

not clear. Our hypothesis was that although frequency might play some role in gender-based categorization of proper names in high-dimensional space, gender information contained in the co-occurrence vectors was also providing categorical information. To limit the possible effects of frequency in gender categorization, another simulation was performed in which the male and female proper names were matched for frequency to control for the greater frequency of male names in the corpus.

Methods

Another set of high-frequency male and female names was taken from the Usenet corpus. However, in this simulation the proper names were chosen in frequency-matched male and female pairs. Co-occurrence vectors were extracted from the model for this list of names and an MDS was performed on these co-occurrence vectors.

Results

Analysis of the MDS demonstrated that male and female names were again discriminated by gender, even when the proper names were matched for frequency (see Figure 2b). An ANOVA provided further inferential support for this finding, showing that male and female intraname frequencies were significantly higher than intername frequencies, $F(1,188) = 4.37, p < .05$; $F(1, 188) = 9.66, p < .01$. These results provide evidence that gender categorization in high-dimensional space is not merely due to the higher frequency of the male names, but rather that gender information is contained in the co-occurrence vectors of the proper names.

Experiment 4c: Gender Information in High Dimensional Proper Name Neighborhoods

Simulation 4b provided evidence that male and female names occupied different locations in high dimensional space, even when the greater frequency of male names was taken into account and controlled for. The purpose of this simulation was to provide further evidence that gender separation in high dimensional space was not an effect of frequency, but rather that useful gender information was contained in the neighborhoods of male and female proper names.

Methods

Neighborhoods consisting of 1000 words were extracted for a list of male and female names. A list of adjectives was then derived from these neighborhoods as seen in Table 3. Adjectives were chosen over other classes of words because they were the least ambiguous, that is, most unlikely to belong to more than one class of words (an example of this is *bill*, which can be a proper name, a non-name noun with several possible meanings, or a verb). Our goal was to determine if gender information would be preserved in these adjectives. Two lists of adjectives, one derived from male names and one derived from female names, were presented to participants, who were asked to judge which list came from the neighborhood of a male name and which list came from the neighborhood of a female name.

Results

Participants were very proficient in determining whether the lists came from male neighborhoods or female neighborhoods, $\chi^2(1, N = 66) = 21.32, p < .001$. These results provide more evidence that gender separation in the high-dimensional name space is not merely frequency-based, but rather that gender relevant information is being preserved in the co-occurrence vectors for proper names.

Table 3: Adjective Neighbors of Male and Female Names

Neighbors of: male names		Neighbors of: female names	
central	sure	sweet	beautiful
special	national	hot	pink
full	double	last	wild
personal	good	new	grand
swedish	digital	lovely	cool

General Discussion

There is a dearth of research into the nature of representations in proper names (Cohen & Burke, 1993). At the same time there is a consensus about several aspects of what the representations must be like; common names lack a substantial semantic component; whereas, famous names have the richer semantic component much like well known objects. We suggest that global co-occurrence models, such as HAL, that develop representations by utilizing the contexts in which names occur are a plausible candidate for

proper name representation. A number of experiments were presented that make this case. First, common and famous proper names occupy their own meaning areas in the high-dimensional space. An analysis of the neighborhoods surrounding common names suggest that there are two factors that may contribute to how retrieval failure is more likely for names. First, common name neighborhoods contain other common names; whereas, neighborhoods of other nouns contain a rich set of cues as to their meaning (see Table 1). Thus, an error in retrieval of a common name results in retrieving the wrong name, but an error in retrieval with other nouns will likely still result in the retrieval of something semantically relevant that may either suffice given context, or could be used to bootstrap the intended concept. This description of the semantic neighborhoods parallels discussion in the proper name literature (Cohen & Burke) about the lack of detailed semantics of common names. The analysis of HAL's name neighborhoods, however, generated a hypothesis that has not been discussed in the literature. We found that the high-dimensional space of the proper name neighborhoods is denser than for frequency-balanced common objects. Thus, any error in retrieval makes an error with common names more likely than with common objects. We think that HAL offers a plausible description (or at least a good starting point) of neighborhoods of famous names. These neighborhoods contain a much richer set of cues, much like that of common objects (see Table 2). Representations such as these provide for a larger set of cues for retrieval and a depth of knowledge such as has been hypothesized by others.

One salient aspect to the proper name representations was gender and this was unexpected. Given the nature of Usenet text, it was thought that the gender distinctions may simply be asymmetries in the frequency of use. However, this was ruled out in Experiment 4b. Lending support to this distinction was Experiment 4c which showed that humans were sensitive to the gender bias of the adjectives in male and female name neighborhoods. Further investigation is required into the contexts in which male and female names occur in order to say more about this result. An important avenue for future research will be to see if these gender cues are good retrieval cues.

Most models of proper names posit a separate component for the storage of proper name lexical/semantic entries (Burton & Bruce, 1993). It is easy to imagine the motivation for this step given the striking dissociations that take place with brain-damaged patients (Semenza, 1995). Furthermore, without a testable representational model it is difficult to imagine proper name information residing with other lexical and semantic information. However, the results presented here suggest that a HAL offers one approach to a single model of meaning in which semantic, grammatical, and common and famous proper name information can coexist in one system. These different kinds of knowledge are in different areas of the high-dimensional space, but we think that considering the possibility of a single meaning

system may facilitate new insights (such as the neighborhood density finding), as well as new questions to ask. Although it may appear surprising that all this information could be encoded in the same meaning system, it is important to bear in mind that the underlying principle in the encoding of conceptual information in HAL is that the contexts in which words appear determines their meaning. The investigation of proper names will likely benefit, like all other areas of cognitive science, by the development of computational models of representation.

Acknowledgements

This research was supported by NSF Presidential Faculty Fellow award SBR-9453406 to Curt Burgess.

References

- Burgess, C. & Lund, K. (1997). Modeling parsing constraints with high-dimensional context space. *Language and Cognitive Processes*, 12, 1-34.
- Burgess, C., Livesay, K., & Lund, K. (1998). Explorations in context space: Words, sentences, discourse. *Discourse Processes*, 25, 211 - 257.
- Burton, A. M., Bruce, V. (1993). Naming faces and naming names: Exploring an interactive activation model of person recognition. *Memory*, 1, 457-480.
- Cohen, G. & Burke, D.M. (1993). Memory for proper names: A review. *Memory*, 1, 249-263
- Lund, K. & Burgess, C. (1996). Producing high-dimensional semantic spaces from lexical co- occurrences. *Behavior Research Methods, Instruments and Computers*, 28, 203-208.
- 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 Bulletin*, 104, 211-240.
- Lund, K., Burgess, C., & Audet, C. (1996). Dissociating semantic and associative word relationships using high-dimensional semantic space. *Proceedings of the Cognitive Science Society* (pp. 603-608). Hillsdale, N.J.: Lawrence Erlbaum Associates.
- McDonald, D.D. (1996). Internal and external evidence in the identification and semantic categorization of proper names. In B. Boguraev & J. Pustejovsky (Eds.) *Corpus Processing for Lexical Acquisition*. MIT Press, Cambridge, Massachusetts.
- Paik, W., Liddy, E.D., Yu, E., & McKenna, M. (1996). Categorizing and standardizing proper nouns for efficient information retrieval. In B. Boguraev & J. Pustejovsky (Eds.) *Corpus Processing for Lexical Acquisition*. MIT Press, Cambridge, Massachusetts.
- Semenza, C. (1995). How names are special: Neuropsychological evidence for dissociable impairment and sparing of proper name knowledge in production. In R. Campbell & M. A. Conway (Eds.) *Broken memories: Case studies in memory impairment*. Blackwell Publishers, Inc, Oxford, England.