ARTICLE IN PRESS

Neuropsychologia ∎ (∎∎∎∎) ∎∎∎–∎∎∎

Contents lists available at ScienceDirect

Neuropsychologia

journal homepage: www.elsevier.com/locate/neuropsychologia

You shall know an object by the company it keeps: An investigation of semantic representations derived from object co-occurrence in visual scenes

Zahra Sadeghi [a,b], James L. McClelland [b], Paul Hoffman [b,c,*]

[a] School of Electrical and Computer Engineering, University of Tehran, Iran
[b] Department of Psychology, Center for Mind, Brain and Computation, Stanford University, Stanford, CA, USA
[c] Neuroscience and Aphasia Research Unit (NARU), School of Psychological Sciences, University of Manchester, Zochonis Building, Oxford Road, Manchester M13 9PL, UK

ARTICLE INFO

Keywords:
Semantic representation
Object knowledge
Latent semantic analysis
Categorisation

ABSTRACT

An influential position in lexical semantics holds that semantic representations for words can be derived through analysis of patterns of lexical co-occurrence in large language corpora. Firth (1957) famously summarised this principle as "you shall know a word by the company it keeps". We explored whether the same principle could be applied to non-verbal patterns of object co-occurrence in natural scenes. We performed latent semantic analysis (LSA) on a set of photographed scenes in which all of the objects present had been manually labelled. This resulted in a representation of objects in a high-dimensional space in which similarity between two objects indicated the degree to which they appeared in similar scenes. These representations revealed similarities among objects belonging to the same taxonomic category (e.g., items of clothing) as well as cross-category associations (e.g., between fruits and kitchen utensils). We also compared representations generated from this scene dataset with two established methods for elucidating semantic representations: (a) a published database of semantic features generated verbally by participants and (b) LSA applied to a linguistic corpus in the usual fashion. Statistical comparisons of the three methods indicated significant association between the structures revealed by each method, with the scene dataset displaying greater convergence with feature-based representations than did LSA applied to linguistic data. The results indicate that information about the conceptual significance of objects can be extracted from their patterns of co-occurrence in natural environments, opening the possibility for such data to be incorporated into existing models of conceptual representation.

© 2014 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/3.0/).

1. Introduction

The structure and content of the conceptual representations of objects are central topics in the study of semantic cognition. It is widely accepted that our understanding of objects and their relationships with one another can be usefully captured by analysing the properties they possess, often referred to as semantic features. A number of large-scale feature listing studies have been conducted, in which participants are asked to generate features for a large set of objects (Cree & McRae, 2003; Devlin, Gonnerman, Andersen, & Seidenberg, 1998; Garrard, Lambon Ralph, Hodges, & Patterson, 2001; Tyler, Moss, Durrant-Peatfield, & Levy, 2000; Vinson, Vigliocco, Cappa, & Siri, 2003; Zannino, Perri, Pasqualetti, Caltagirone, & Carlesimo, 2006). In such studies, participants tend to produce features derived from perceptual experience (e.g., lemons are yellow), functional features concerned with behaviours or goals associated with the object (lemons are used to make drinks) and more abstract information that can typically only be expressed verbally (lemons are a type of citrus fruit). On this view, two objects are conceptually related to the extent that they share similar features; so oranges are semantically linked with lemons because they too are citrus fruits and are used to make drinks. Feature generation studies of this kind have strongly endorsed the view that object knowledge is organised in terms of taxonomic category. Objects that belong to the same taxonomic category tend to share features (Cree & McRae, 2003) and, moreover, items that share many features with other items from their category are judged to be more prototypical members of the category (Garrard et al., 2001). Dilkina and Lambon Ralph (2012) recently demonstrated that items within the same category

* Corresponding author at: Neuroscience and Aphasia Research Unit (NARU), School of Psychological Sciences, Zochonis Building, University of Manchester, Oxford Road, Manchester M13 9PL, UK. Tel.: +44 161 275 7338; fax: +44 161 275 2873.
E-mail address: paul.hoffman@manchester.ac.uk (P. Hoffman).

http://dx.doi.org/10.1016/j.neuropsychologia.2014.08.031
0028-3932/© 2014 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/3.0/).

most frequently shared features that referred to their perceptual qualities, though functional and more abstract encyclopaedic features were also somewhat linked to taxonomic organisation. The patterning of correlations amongst features and the relative salience of different types of feature have also been shown to vary across living and non-living things (Farah & McClelland, 1991; Garrard et al., 2001; Tyler et al., 2000). Living things are more strongly associated with perceptual features, for example, and manufactured artefacts with functional features. These differences have been proposed to account for patterns of category-selective semantic deficits sometimes observed in a variety of neurological conditions (Cree & McRae, 2003; Farah & McClelland, 1991; Warrington & Shallice, 1984).

The feature-based approach to object knowledge has proved fruitful, with a number of models of object knowledge assuming that object concepts are structured in terms of their featural similarity (Collins & Quillian, 1969; McRae, deSa, & Seidenberg, 1997; Rogers et al., 2004; Rogers & McClelland, 2004; Tyler et al., 2000; Vigliocco, Vinson, Lewis, & Garrett, 2004). The idea that taxonomic category is a key organising principle for object concepts has also guided recent neuroimaging studies that have used multi-voxel pattern analysis to investigate representational structure (Devereux, Clarke, Marouchos, & Tyler, 2013; Fairhall & Caramazza, 2013; Kriegeskorte et al., 2008; Peelen & Caramazza, 2012). Some limitations of the feature-based approach have been noted, however. It has been suggested that the feature generation task is biased towards features that distinguish objects from their category neighbours and towards aspects of information that can be easily expressed verbally (Hoffman & Lambon Ralph, 2013; Rogers et al., 2004). Another, perhaps more fundamental, limitation is the fact that participants generating semantic features are asked to consider each object in isolation. The relationships between objects are therefore inferred indirectly, in terms of their feature overlap. This is not representative of our natural experience with objects. Environments typically contain many objects and most activities require us to interact with multiple objects simultaneously, which often have few features in common. To extend our earlier example, in order to make lemonade, life must give you not only lemons but water, sugar and a jug. How does the co-occurrence of these objects influence our conceptual representations of each of them?

An alternative approach to semantic representation has developed in the field of computational linguistics, based on the idea that semantic representations of words can be derived through statistical analysis of their distribution in large text corpora (Firth, 1957; Griffiths, Steyvers, & Tenenbaum, 2007; Landauer & Dumais, 1997; Lund & Burgess, 1996; Rohde, Gonnerman, & Plaut, 2006). The central tenet underpinning the distributional approach is the idea that words that occur in similar linguistic contexts are related in meaning. On this view, oranges and lemons would be considered similar because they co-occur with a similar set of words in natural language. For example, we might expect both *orange* and *lemon* to frequently occur in sentences that contain words like *squeeze, cut, peel, pips, juice* and *marmalade*. On the face of it, this does not sound so different to the featural approach. However, the distributional approach allows for the possibility that objects from different taxonomic categories which share few features may nevertheless share a semantic relationship (e.g., *lemon* and *ice* may be considered semantically related because both words are used when we talk about making drinks). These associative or thematic relationships are known to play an important role in lexical-semantic processing. For example, significant semantic priming effects occur for word pairs that share an associative relationship as well as items that share semantic features (Alario, Segui, & Ferrand, 2000; Perea & Gotor, 1997; Seidenberg, Waters, Sanders, & Langer, 1984). Furthermore, children readily group objects according to their associative relationships and may even prefer this to grouping by taxonomic similarity (Kagan, Moss, & Sigel, 1963; Smiley & Brown, 1979), suggesting that associations play an important role in the development of concepts. Therefore lexical co-occurrence likely serves as an additional source of constraint over the structuring of object concepts, since it is able to capture associative relationships between items that share few features. However, semantic models based on the distributional principle have been criticised because they rely solely on linguistic data and therefore do not take into account, at least in any direct way, the sensory-motor information available when we perceive and interact with objects in the real world (Andrews, Vigliocco, & Vinson, 2009; Glenberg & Robertson, 2000). Linguistic corpora may code perceptual experiences indirectly, of course, through verbal descriptions of sensory experiences.

Feature lists and lexical co-occurrence provide two differing perspectives on the conceptual relationships among objects. There is now evidence that true semantic representation requires a combination of these two sources of data. In an innovative study, Andrews et al. (2009) used a Bayesian probabilistic model to generate semantic representations for objects based jointly on feature lists and word co-occurrence information obtained from a text corpus. The resultant representations provided a better fit to a range of empirical data than those derived from either data source in isolation. This suggests that our understanding of the relationships between objects is based partly on shared properties and partly on knowledge of their co-occurrence. Other researchers have used related statistical methods to integrate feature knowledge with data about concept co-occurrence (Durda, Buchanan, & Caron, 2009; Johns & Jones, 2012; Steyvers, 2010). All of these studies have used linguistic corpus data as the basis for inferring patterns of contextual co-occurrence among objects. However, much of our experience of concrete objects is non-verbal: in addition to using words that refer to objects together in sentences, we also perceive combinations of objects directly in different environments. For example, we frequently see oranges and lemons together in fruit bowls. This direct experience of object co-occurrence potentially provides a rich additional source of information about object concepts, beyond that provided by feature lists and lexical co-occurrence; however, its potential contribution to semantic knowledge has not been assessed. In this study, we investigated whether meaningful semantic information can be derived from patterns of object co-occurrence, by applying latent semantic analysis (LSA) to a set of labelled photographs that depict collections of objects in a variety of natural scenes (see Fig. 1 for examples). LSA is commonly used to derive high-dimensional semantic representations for words based on underlying similarities in the verbal contexts in which they are used (Landauer & Dumais, 1997). Here, we used the same technique to derive high-dimensional semantic representations for objects based on underlying similarities in the environments in which they appear. We compared semantic representations derived in this way with (a) representations based on feature lists (McRae, Cree, Seidenberg, & McNorgan, 2005) and (b) representations obtained through the traditional application of LSA to linguistic corpus data. We aimed to explore the degree to which information derived from environmental co-occurrence provided similar or complementary information about objects as these other two sources.

2. Method

2.1. Processing of the scene dataset

We used latent semantic analysis (LSA; Landauer & Dumais, 1997) to investigate patterns in visual object co-occurrence. LSA is a well-known technique for constructing semantic representations based on lexical co-occurrence in text

Download English Version:

https://daneshyari.com/en/article/7319701

Download Persian Version:

https://daneshyari.com/article/7319701

Daneshyari.com