Title
All parts are not created equal: SIAM-LSA

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 26(26)

ISSN
1069-7977

Author
Wiemer-Hastings, Peter

Publication Date
2004

Peer reviewed

All parts are not created equal: SIAM-LSA

Peter Wiemer-Hastings
peterwh@cti.depaul.edu
DePaul University
School of Computer Science,
Telecommunications, and Information Systems
243 S. Wabash
Chicago IL 60604

The ability to assess the similarity of objects in the world is fundamentally important to our survival. Many theories have been proposed for modeling human similarity judgments. Most of these theories involve comparing the sets of features of the compared items to determine the overlap between them. Many of them completely ignore the structure of the objects and the relationships between the parts. Goldstone (1994) showed that such systems fail to account for human similarity ratings of structured data. His SIAM system used a (non-learning) connectionist architecture to create correspondences between objects and their features in different scenes. Excitatory connections reinforced coherent mappings between objects (e.g. ObjectA to ObjectC and ObjectB to ObjectD). Inhibitory connections fought against redundant or contradictory mappings. Likewise, connections between the features of objects either supported or inhibited each other and the corresponding object–object connections. SIAM's connectionist architecture allowed it to take into account the structure of the scenes and the objects as well as the similarity of the features.

Goldstone examined similarity ratings of visual scenes. His approach represented a scene as a spatially related set of objects (for example, pairs of schematic butterflies). Each object has a set of parts each of which has some value. For example, one of Goldstone's butterflies could be represented as: `(object1 (head square) (tail zig-zag) (body-shading white) (wing-shading checkered))`.

In previous research, we have explored the use of Latent Semantic Analysis (LSA) for judging the semantic similarity of a given sentence to a set of alternative target sentences. Although LSA has been shown to match the reliability of raters with intermediate domain knowledge, the correlation between LSA and human ratings is still somewhat disappointing, generally below 0.5 in a number of studies (Wiemer-Hastings, Wiemer-Hastings, & Graesser, 1999). In recent research, we have pursued the general hypothesis that including structural knowledge would improve the correspondence between human and LSA ratings. We found that by performing syntactic analysis of the source and target sentences and separately comparing their subjects, objects, and verbs with LSA, we could reduce the error by over 10% (Wiemer-Hastings & Zipitria, 2001).

In the current research, we explored the use of SIAM to combine the analysis of the structural aspects of the sentences with the semantic similarity ratings provided by LSA. To map this approach to sentences, we broke the inputs into subject, verb, object, and indirect object parts. Thus, a simple representation of the sentence "The dog bit a man" as an object would be: `(object1 (verb "bit") (subject "The dog") (object "a man"))`. The advantage of SIAM-LSA over the previous model (Structured LSA, or SLSA) is that its connectionist architecture allows the different components to "compete" for correspondence, instead of relying on a direct mapping of subject, verb, and object segments. Our basic hypothesis was that SIAM-LSA would provide a closer match to human ratings than SLSA. A secondary hypothesis was that providing a salience value to give differential weight to the different structural components of the sentences would better match human ratings.

In our experiment, we compared human ratings with the basic SIAM-LSA system and the system augmented with salience values. Our results did not support the basic hypothesis. In fact, SIAM-LSA performed worse than LSA or SLSA. When we included empirically derived weights which accentuated verb and object matches but completely devalued subject matches, the ratings correlated with human ratings $r = 0.59$, another 10% reduction in the error over SLSA. In accordance with (Resnik, 1993), this suggests that humans essentially ignore the role of syntactic subjects when matching sentence meanings.

References

Goldstone, R. (1994). Similarity, Interactive Activation, and Mapping. *Journal of Experimental Psychology*, *20*(1), 3–28.

Resnik, P. (1993). *Selection and Information: A class-based approach to lexical relationships*. Ph.D. thesis, University of Pennsylvania, Philadelphia, PA.

Wiemer-Hastings, P., Wiemer-Hastings, K., & Graesser, A. (1999). How Latent is Latent Semantic Analysis?. In *Proceedings of the Sixteenth International Joint Congress on Artificial Intelligence*, pp. 932–937 San Francisco. Morgan Kaufmann.

Wiemer-Hastings, P., & Zipitria, I. (2001). Rules for Syntax, Vectors for Semantics. In *Proceedings of the 23^{rd} Annual Conference of the Cognitive Science Society* Mahwah, NJ. Erlbaum.