A Neurobiologically Motivated Analysis of Distributional Semantic Models

Akira Utsumi (utsumi@uec.ac.jp)
Department of Informatics, The University of Electro-Communications
1-5-1, Chofugaoka, Chofushi, Tokyo 182-8585, Japan

Abstract

The pervasive use of distributional semantic models or word embeddings is due to their remarkable ability to represent the meanings of words for both practical application and cognitive modeling. However, little has been known about what kind of information is encoded in text-based word vectors. This lack of understanding is particularly problematic when distributional semantics is regarded as a model of semantic representation for abstract concepts. This paper attempts to reveal the internal knowledge encoded in distributional word vectors by the analysis using Binder et al.'s (2016) brain-based vectors, explicitly structured conceptual representations based on neurobiologically motivated attributes. In the analysis, the mapping from text-based vectors to brain-based vectors is trained and prediction performance is evaluated by comparing the estimated and original brain-based vectors. The analysis demonstrates that social and cognitive information is predicted with the highest accuracy by text-based vectors, but emotional information is not predicted so accurately. This result is discussed in terms of embodied theories for abstract concepts.

Keywords: Distributional semantic models; Word vectors; Brain-based representation; Embodied cognition; Emotional and social information; Abstract concepts

Introduction

One of the most important advances in the study of semantic processing is the development of distributional semantic models for representing word meanings. In the distributional semantic model, words are represented as high-dimensional vectors, which can be learned from the distributional statistics of word occurrence in large collections of text. Any words that occur in the corpus can be learned regardless of their part-of-speech class, abstractness, novelty and familiarity. This is an important advantage of text-based distributional semantic models over other spatial models of semantic representation such as feature-based (Andrews, Vigliocco, & Vinson, 2009) and image-based vectors (Silberer, Ferrari, & Lapata, 2017).

Word vectors have been employed in a variety of research fields and many successful results have been obtained. In the field of natural language processing (NLP), deep learning has recently been applied to a number of NLP tasks such as machine translation and automatic summarization, and achieved the impressive performance as compared to the traditional statistical methods. One of the reasons for the successful results is the use of word vectors as semantic representations for the input and output of recurrent neural networks (Goldberg, 2017). Research on cognitive science also benefits greatly from distributional semantic models (Jones, Willits, & Dennis, 2015). Word vectors have been demonstrated to explain a number of cognitive phenomena relevant to semantic memory or mental lexicon, such as word association (Jones, Gruenenfelder, & Recchia, 2017; Utsumi, 2015), semantic priming (Mandera, Keuleers, & Brysbaert, 2017), semantic transparency (Marelli & Baroni, 2015) and conceptual combination (Vecchi, Marelli, Zamparelli, & Baroni, 2017). Furthermore, recent brain imaging studies have demonstrated that distributional word vectors have a powerful ability to predict the neural brain activity in cerebral cortex evoked by lexical processing (Mitchell et al., 2008; Huth, de Heer, Griffiths, Theunissen, & Gallant, 2016; Anderson, Kiela, Clark, & Poesio, 2017). These findings show that distributional semantic models reflect the representational structure of semantic knowledge in the brain.

Despite the fact that successful results are obtained in many research fields, relatively little has been known about what kind of information or knowledge is encoded in word vectors. Some existing studies have addressed this question, demonstrating that text-based word vectors reflect perceptual (Louwerse & Connell, 2011; Riordan & Jones, 2011), emotional (Recchia & Louwerse, 2015; Tillmand & Louwerse, 2018), and social (Hutchinson & Louwerse, 2018) information. However, no direct comparison among a wide variety of information has been made with respect to the representational ability of distributional semantic models. This lack of understanding makes distributional semantic models unable to predict human language behavior and performance at the same level of detail and precision of other cognitive models. It also limits further improvements on the practical performance of word vectors for many NLP tasks.

In this paper, therefore, we attempt to reveal the internal knowledge encoded in text-based word vectors by comprehensively exploring their representational ability of various types of information. Our approach to this problem is to simulate a brain-based semantic representation (Binder et al., 2016) using text-based vectors. This semantic representation comprises 65 attributes (listed in Table 1) based entirely on functional divisions in the human brain. Each word is represented as a 65-dimensional vector and each dimension represents the salience of the corresponding attribute, namely the degree to which the concept referred to by that word is related to that attribute. Because these attributes are based on not only sensorimotor experiences but also affective, social, and cognitive experiences, we can analyze distributional word vectors considering a wide variety of information. In the analysis, we trained the mapping from the text-based vectors to the brain-based vectors, by which brain-based vectors of untrained words are predicted.

The secondary purpose of this paper is to discuss the relationship between the embodied theory for abstract concepts and distributional semantic models from the results of the analysis. Recently, it has been accepted that language or linguistic experience is much more important for representing and acquiring abstract concepts than for concrete concepts, because abstract concepts are unlikely to be grounded in perceptual and sensorimotor experiences (Borghi et al., 2017).[1] A number of approaches have been proposed to explain the

Table 1: 65 attributes used in brain-based vectors

Domain	Attributes
Vision	Vision, Bright, Dark, Color, Pattern, Large, Small, Motion, Biomotion, Fast, Slow, Shape, Complexity, Face, Body
Somatic	Touch, Temperature, Texture, Weight, Pain
Audition	Audition, Loud, Low, High, Sound, Music, Speech
Gustation	Taste
Olfaction	Smell
Motor	Head, UpperLimb, LowerLimb, Practice
Spatial	Landmark, Path, Scene, Near, Toward, Away, Number
Temporal	Time, Duration, Long, Short
Causal	Caused, Consequential
Social	Social, Human, Communication, Self
Cognition	Cognition
Emotion	Benefit, Harm, Pleasant, Unpleasant, Happy, Sad, Angry, Disgusted, Fearful, Surprised
Drive	Drive, Needs
Attention	Attention, Arousal

Table 2: Example of words represented as brain-based vectors

Category	Word	Category	Word
plant	apricot, rose, tree	human	actor, girl, parent
vehicle	car, subway, boat	social action	celebrate, help
place	airport, lake, lab	visual property	black, new, dark

role of language as a simple shortcut (Barsalou, Santos, Simmons, & Wilson, 2008) or indirect grounding in perceptual or sensorimotor experiences (Louwerse, 2011; Dove, 2014), and the need for other information such as emotional (Kousta, Vigliocco, Vinson, Andrews, & Del Campo, 2011) and social information (Borghi & Binkofski, 2014). Because text-based word vectors can be regarded as realizations of linguistic experiences, the analysis of internal knowledge encoded in text-based word vectors is expected to provide implications for recent embodied approaches to abstract concepts.

Method

In order to explore the information encoded in distributional word vectors, we evaluated how accurately they can simulate Binder et al.'s (2016) brain-based vectors. The simulation was performed by training the mapping from text-based vectors to brain-based vectors and applying the trained mapping to the text-based vectors of untrained words. Prediction performance was evaluated by comparing the estimated brain-based vectors with the original brain-based vectors.

Brain-based Vectors

As mentioned above, we used Binder et al.'s (2016) brain-based componential representation of words as a gold standard. They provided 65-dimensional vectors of 535 words comprising 434 nouns, 62 verbs and 39 adjectives, some of which are listed in Table 2. The 65 dimensions listed in Table 1 correspond to neurobiologically plausible attributes whose neural correlates have been well described. These attributes were selected according to two fundamental principles; they correspond to distinguishable neural processors that can be identified by an extensive body of evidence from brain imaging and neurological studies, and they can contribute to concept acquisition and composition.

Elements of the brain-based vector represent the degree of salience of attributes for the target word. Binder et al. (2016) collected these values using Amazon Mechanical Turk. Participants of the survey were given a single word and questions such as "To what degree do you think of this thing as a characteristic or defining color " (for the attribute *Color*) with some examples, and asked to rate the degree on a 7-point scale ranging from 0 to 6. Collected ratings were averaged for each word and attributed after data screening, and these mean ratings were used in brain-based vectors.

Word Vectors

In order to ensure the generality of the findings obtained through the analysis, we constructed six semantic spaces, which were obtained from the combinations of three distributional semantic models (SGNS, GloVe, PPMI) and two corpora (COCA and Wikipedia). As distributional semantic models, we used three representative models, namely skip-gram with negative sampling (SGNS; Mikolov, Chen, Corrado, & Dean, 2013), GloVe (Pennington, Socher, & Manning, 2015) and positive pointwise mutual information (PPMI) with SVD (Bullinaria & Levy, 2007). SGNS and GloVe are prediction-based models that train word vectors by predicting context words on either side of a target word, while PPMI is a counting-based model that trains word vectors by counting and weighting word occurrences. We set a vector dimension $d = 300$ and a window size $w = 10$ for all semantic spaces.

Two corpora used in the analysis were English Wikipedia dump of enwiki-20160601 (Wiki) and Corpus of Contemporary American English (COCA). The Wiki and COCA corpora include 1.89G and 0.56G word tokens, respectively. We built a vocabulary from frequent words that occur 50 times or more in Wiki corpus [2], or 30 times or more in COCA corpus. As a result, the vocabulary of Wiki and COCA contained 291,769 and 108,230 words, respectively. These two corpora differ in that Wiki is a raw text corpus that is untagged and unlemmatized, while COCA is a fully tagged and lemmatized corpus. For Wiki corpus, raw texts were extracted from the dump files using `WikiExtractor.py` [3] and no other preprocessing such as lemmatization was applied.

Learning Methods for the Mapping

We used two learning methods, namely linear transformation (LT) and multi-layer perceptron (MLP). LT trains a mapping matrix $\boldsymbol{M}$ such that $\boldsymbol{B} = \boldsymbol{W}\boldsymbol{M}$ where $\boldsymbol{B}$ is the matrix with brain-based word vectors as rows and $\boldsymbol{W}$ is a matrix with text-based word vectors as rows. MLP trains a neural network with one

[1] Note that there are some suggestions that some abstract concepts are grounded in sensorimotor experiences (Connell & Lynott, 2012; Dreyer & Pulvermüller, 2018).

[2] Out of 535 words for brain-based vectors, only one word "joviality" was not selected as frequent words for Wiki corpus. Hence, we added it to the vocabulary for Wiki corpus.

[3] `http://medialab.di.unipi.it/wiki/Wikipedia_Extractor`

Table 3: Mean correlations over all attributes

		SGNS	GloVe	PPMI
Wikipedia	MLP	0.576	0.522	0.483
	LT	0.549	0.450	0.429
COCA	MLP	0.634	0.554	0.440
	LT	0.598	0.494	0.454

hidden layer comprising 150 sigmoid units and a linear output layer. In both methods, the mapping was trained by minimizing the mean squared error, and gradient descent with AdaGrad was used as an optimization method.

Estimation of brain-based vectors from text-based vectors was performed by a leave-one-out cross validation procedure. For each of the 535 words, we trained the mapping between brain-based and text-based vectors of the remaining 534 words and estimated a brain-based vector for the target word using the trained mapping. By repeating this procedure for all words as a target, we obtained $\widehat{\boldsymbol{B}}$ with estimated brain-based vectors as rows.

Performance Measure

Prediction performance of the estimated vectors was measured using Spearman's rank correlation ρ between the estimated brain-based matrix $\widehat{\boldsymbol{B}}$ and the original matrix $\boldsymbol{B}$.[4] We performed two analyses: column-wise and row-wise matrix correlation. The column-wise matrix correlation indicates the estimation accuracy for each attribute, while the row-wise correlation indicates the accuracy for each word.

Result

Correlation Analysis by Attribute

We evaluated the prediction accuracy for attributes by computing column-wise matrix correlations between the estimated and original brain-based vector spaces. Figure 1 shows correlation coefficients for 65 attributes. In addition, these results are summarized in Figure 2, which depicts mean correlations averaged over attributes of the same domain.

Although in this paper we are not concerned with the performance difference between word vectors, Table 3 shows that SGNS achieved the best prediction performance in the three models, and word vectors trained using the COCA corpus were superior to those of the Wiki corpus. In addition, as expected, MLP trained better mappings than LT. A three-way ANOVA on Fisher's z-transformed correlations revealed that all these differences were significant, $F(2,128)=261.3$, $p<.001$ for model; $F(1,64)=66.8$, $p<.001$ for corpus; and $F(1,64)=186.4$, $p<.001$ for learning method.

Despite these differences of overall performance, Figures 1 and 2 demonstrate that relative performance among attributes did not differ, regardless of distributional semantic model, corpus and training method. To confirm this statistically, we

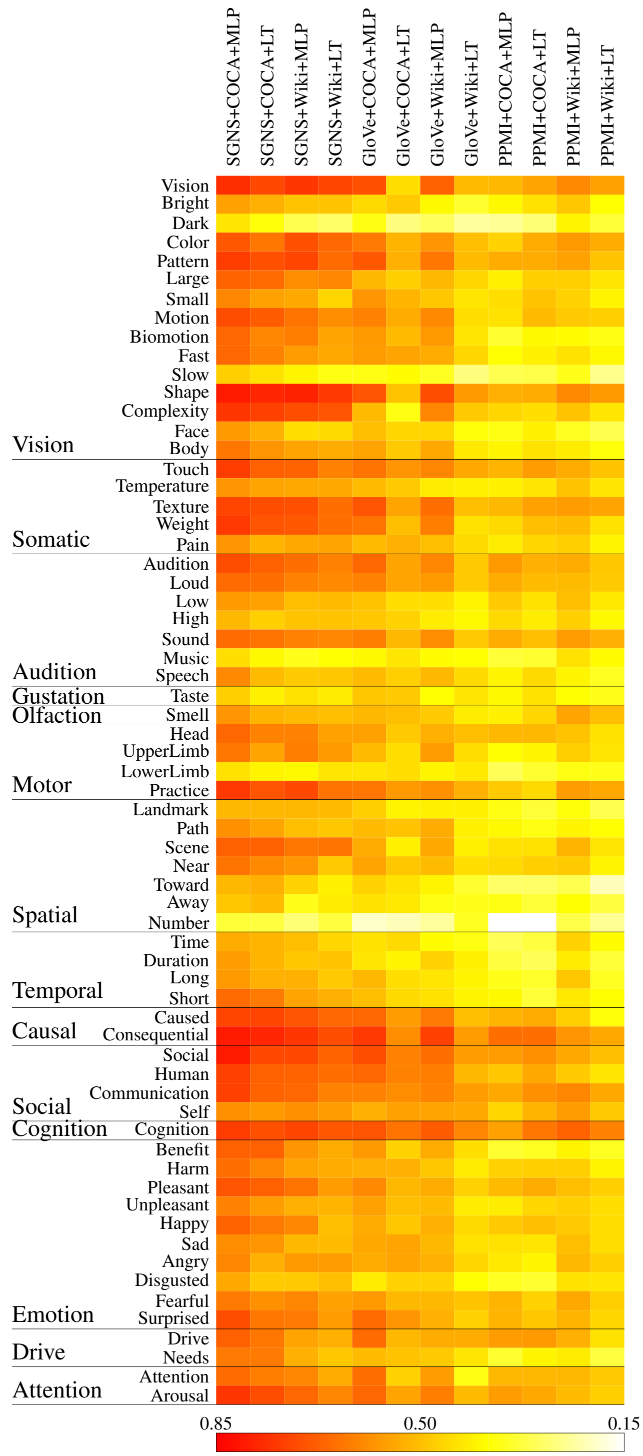

Figure 1: Correlations between the estimated and original brain-based vectors for 65 attributes. Each row corresponds to the results of an attribute and each column shows the results of combinations of models (SGNS, GloVe, PPMI), corpora (COCA, Wiki) and training methods (MLP, LT).

computed Spearman's correlations of 65 attribute correlations for all pairs of 12 different results. Correlations of correlations (ranging from 0.58 to 0.97) were all statistically significant ($p<.01$, false discovery rate corrected).

Attributes in *Causal*, *Cognition*, *Social*, and *Attention* do-

[4] Mean squared error can also be a measure for prediction performance. However, we are interested in the similarity of order, rather than of absolute value, between the original and estimated vectors, and thus we used rank correlations in this paper.

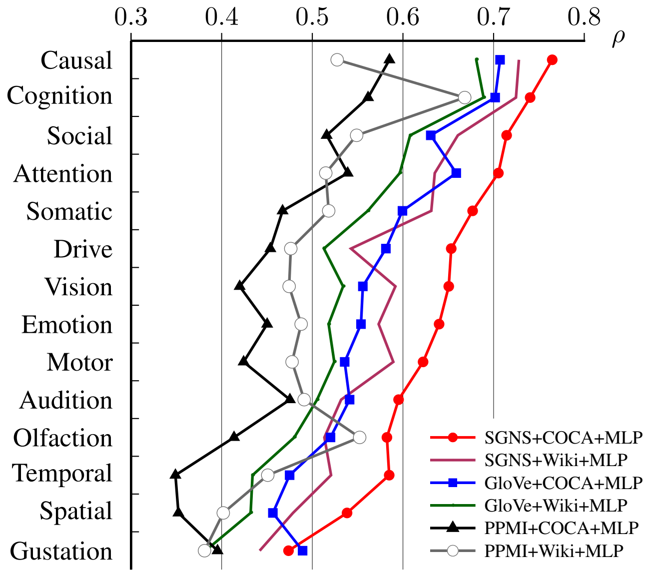

Figure 2: Mean correlations per attribute domain. Only the results for MLP are shown for simplicity.

mains were generally predicted with higher accuracy; their correlations of SGNS+COCA+MLP exceeded 0.7. In particular, *Causal* and *Cognition* domains achieved significantly higher correlations than all other domains.[5] The correlations of *Social* and *Attention* domains were also significantly higher than those of nine domains (from *Drive* to *Gustation* in Figure 2). In other words, the information of these attributes, which primarily characterize abstract concepts (Binder et al., 2016), is likely to be encoded in text-based word vectors. This finding seems to suggest that abstract concepts can be largely acquired through linguistic experiences.

Although more difficult to predict than these attributes, perceptual attributes in *Vision*, *Somatic*, *Audition*, and *Olfaction* and motor attribute in *Motor* can be somewhat predicted from text-based word vectors. Correlations of all these domains were significantly higher than spatiotemporal domains *Temporal* and *Spatial* and one perceptual domain *Gustation*. Furthermore, Figure 1 shows that some sensorimotor attributes such as *Vision*, *Pattern*, *Shape*, *Texture* and *Practice* were predicted as accurately as abstract attributes. These findings suggest that text-based word vectors can encode some kinds of sensorimotor information; this is consistent with some existing findings (e.g., Louwerse & Connell, 2011). On the other hand, spatiotemporal attributes in *Temporal* and *Spatial* were most difficult to predict from text-based vectors. This supports the embodied view that spatial information is heavily grounded in perceptual experiences (Zwaan & Yaxley, 2003).

A somewhat surprising result was that emotional attributes were not predicted as accurately as social and cognitive ones, although a large number of NLP studies have demonstrated

[5]After confirming that a one-way ANOVA on Fisher's z-transformed correlations showed a significant main effect of attribute, $F(13, 143) = 60.5$, $p < .001$, we assessed the statistical significance of pairwise comparisons using Bonferroni adjustment ($p < .05$).

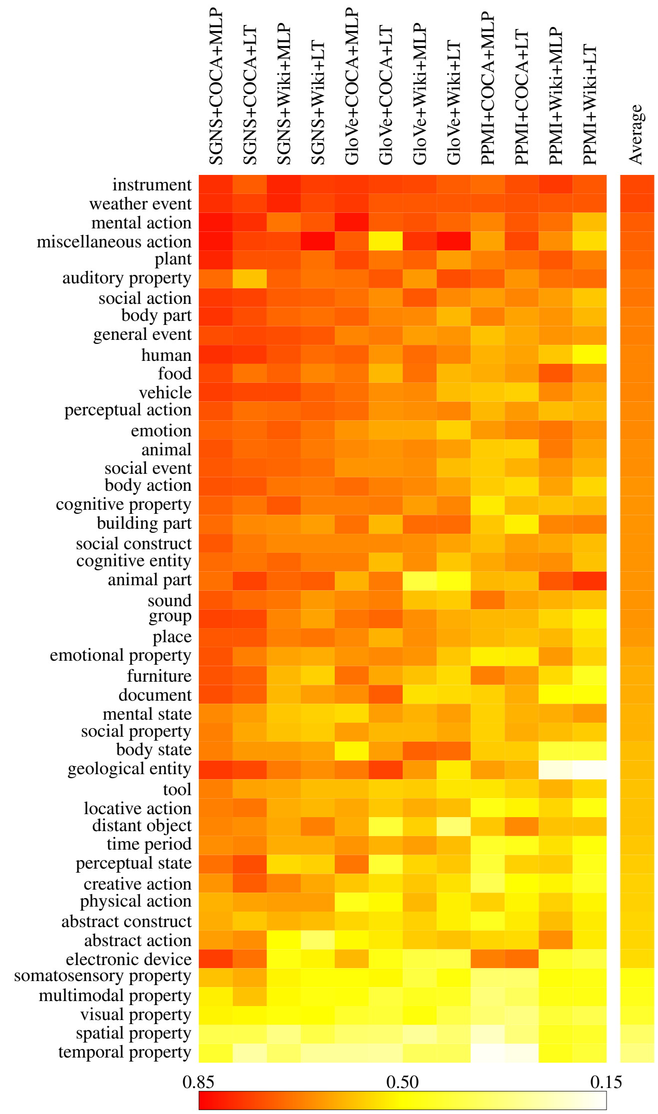

Figure 3: Mean correlations between the estimated and original brain-based vectors for 47 word categories. Each row corresponds to the results of a word category.

successful results of sentiment analysis (Taboada, 2016) and emotional judgment (Recchia & Louwerse, 2015). The domain *Emotion* showed significantly higher correlations than only the domains *Temporal*, *Spatial* and *Gustation*. This result implies that emotional information is more likely to be acquired from direct emotional experiences than from linguistic ones. It is consistent with the recent embodied view that emotional experiences are required for grounding abstract concepts (Kousta et al., 2011; Vigliocco et al., 2014).

Correlation Analysis by Word

We computed row-wise correlations between the estimated and original brain-based matrices, and then averaged these 535 correlations according to 47 word categories. These word categories are provided a priori by Binder et al. (2016) and reflect grammatical classes (i.e., noun, verb, adjective) and semantic classes.[6] Figure 3 shows mean correlations per word category. As in the case of the attribute analysis, relative performance differences of word categories were similar among semantic spaces and training methods. Spearman's correlations of 47 category correlations (ranging from 0.27 to 0.92) were almost significant ($p < .05$, false discovery rate corrected), although three out of 66 correlations of correlations were not significant.

The overall result was that brain-based vectors for human-related categories such as *mental action*, *social action*, *human* and *social event* were relatively better predicted from text-based word vectors. Emotional and cognitive categories such as *emotion* and *cognitive property* were predicted well, but with lower accuracy than human-related categories. These results are consistent with the findings obtained by the attribute analysis. On the other hand, other abstract concepts, in particular many categories of action and property, were difficult to predict from text-based word vectors. One possible reason for this result may be that the number of verbs and adjectives in the vocabulary is much smaller as compared to nouns, and thus verbs and adjectives are difficult to train. It can also be interpreted as suggesting that distributional semantic models may be insufficient for representing some kinds of abstract concepts, and other experiences than linguistic one would be required (e.g., Borghi et al., 2017).

Interestingly, many artifact categories such as *instruments*, *food*, and *vehicle*, and some natural objects such as *plant* and *animal* showed higher prediction performance. There is no doubt that, as the embodied theory of language argues, these concrete words or concepts are grounded in perceptual and sensorimotor experiences, but some kinds of concrete concepts, in particular artifacts, may be able to be represented (or indirectly grounded) by text-based word vectors.

Discussion

In this paper, we have demonstrated that text-based distributional word vectors can predict social and cognitive information quite accurately, but the accuracy of emotional information is not so high. Given the existing empirical findings on the importance of emotion for abstract concepts (Vigliocco et al., 2014; Buccino, Colagè, Gobbi, & Bonaccorso, 2016), this result suggests that direct emotional experiences are necessary for grounding abstract concepts, and thus lends support to some embodied theories (Kousta et al., 2011; Vigliocco et al., 2014). On the other hand, some other embodied theories such as WAT theory (Borghi & Binkofski, 2014) have argued that social experiences also play an important role in representation of abstract concepts. However, the result of our analysis that social information can be conveyed by language may diminish the importance of social experiences for abstract concepts. Note also that the need of social-cognitive ability is not specific to abstract concepts; concrete concepts are acquired and processed through social abilities such as a Theory of Mind (e.g., Bloom, 2000).

It was also found from the analysis that perceptual, sensorimotor and spatiotemporal information is relatively less likely to be encoded in word vectors. This difficulty often leads to the criticism that distributional semantic models are inadequate models of semantic representation (Glenberg & Robertson, 2000). This result is also consistent with the findings of multimodal distributional semantics that inclusion of visual information improves semantic representation for concrete words (e.g., Kiela, Hill, Korhonen, & Clark, 2014).

Nevertheless, the analysis also showed the possibility that some perceptual information and representation of some concrete concepts can be derived from distributional linguistic statistics, as already demonstrated by other studies (Louwerse & Connell, 2011; Riordan & Jones, 2011). This possibility suggests that the role of language in semantic representation of concrete concepts is more important than what the embodied theories of meaning have expected.

Of course, the analysis presented in this paper has some limitations. One important limitation is that the brain-based vectors represent the salience of attributes that characterize concepts, but do not necessarily represent the value of salient attributes. For some attributes such as *Bright* and *Happy*, their value is indistinguishable from their salience, but many other attributes such as *Color* and *Human* have distinct values independent of their salience. Hence, the analysis in this paper cannot reveal the representational power of attribute values. Our analysis is also limited within a small set of vocabulary words. To generalize and refine the findings presented in this paper, we have to evaluate a much larger set of vocabulary words that are not included in Binder et al.'s (2016) dataset. Further research is needed to overcome these limitations.

[6]Note that word categories provided online slightly differ from those listed in Binder et al.'s (2016) article. In this paper, we used the online version of word categories.

Acknowledgments

This research was supported by JSPS KAKENHI Grant Numbers JP15H02713 and SCAT Research Grant.

References

Anderson, A. J., Kiela, D., Clark, S., & Poesio, M. (2017). Visually grounded and textual semantic models differentially decode brain activity associated with concrete and abstract nouns. *Transactions of the Association for Computational Linguistics*, *5*, 17–30.

Andrews, M., Vigliocco, G., & Vinson, D. (2009). Integrating experiential and distributional data to learn semantic representations. *Psychological Review*, *116*, 463–498.

Barsalou, L. W., Santos, A., Simmons, W. K., & Wilson, C. D. (2008). Language and simulation in conceptual processing. In M. de Vega, A. Glenberg, & A. Graesser (Eds.), *Symbols and embodiment: Debates on meaning and cognition* (pp. 245–283). New York: Oxford University Press.

Binder, J. R., Conant, L. L., Humphries, C. J., Fernandino, L., Simons, S. B., Aguilar, M., & Desai, R. H. (2016). Toward a brain-based componential semantic representation. *Cognitive Neuropsychology*, *33*(3–4), 130–174.

Bloom, P. (2000). *How children learn the meanings of words*. MIT Press.

Borghi, A. M., & Binkofski, F. (2014). *Words as social tools: An embodied view on abstract concepts*. New York: Springer.

Borghi, A. M., Binkofski, F., Castelfranchi, C., Cimatti, F., Scorolli, C., & Tummolini, L. (2017). The challenge of abstract concepts. *Psychological Bulletin*, *143*, 263–292.

Buccino, G., Colagè, I., Gobbi, N., & Bonaccorso, G. (2016). Grounding meaning in experience: A broad perspective on embodied language. *Neuroscience & Biobehavioral Reviews*, *69*, 69–78.

Bullinaria, J. A., & Levy, J. P. (2007). Extracting semantic representations from word co-occurrence statistics: A computational study. *Behavior Research Methods*, *39*(3), 510–526.

Connell, L., & Lynott, D. (2012). Strength of perceptual experience predicts word processing performance better than concreteness or imageability. *Cognition*, *125*(3), 452–465.

Dove, G. (2014). Thinking in words: Language as an embodied medium of thought. *Topics in Cognitive Science*, *6*, 371–389.

Dreyer, F. R., & Pulvermüller, F. (2018). Abstract semantics in the motor system? — An event-related fMRI study on passive reading of semantic word categories carrying abstract emotional and mental meaning. *Cortex*, *100*, 52–70.

Glenberg, A., & Robertson, D. (2000). Symbol grounding and meaning: A comparison of high-dimensional and embodied theories of meaning. *Journal of Memory and Language*, *43*, 379–401.

Goldberg, Y. (2017). *Neural network methods for natural language processing*. Morgan & Claypool Publishers.

Hutchinson, S., & Louwerse, M. (2018). Extracting social networks from language statistics. *Discourse Processes*, in press.

Huth, A. G., de Heer, W. A., Griffiths, T. L., Theunissen, F. E., & Gallant, J. L. (2016). Natural speech reveals the semantic maps that tile human cerebral cortex. *Nature*, *532*, 453–458.

Jones, M. N., Gruenenfelder, T. M., & Recchia, G. (2017). In defense of spatial models of semantic representation. *New Ideas in Psychology*, in press.

Jones, M. N., Willits, J., & Dennis, S. (2015). Models of semantic memory. In J. R. Busemeyer, Z. Wang, J. T. Townsend, & A. Eidels (Eds.), *Oxford handbook of mathematical and computational psychology* (pp. 232–254). New York, NY: Oxford University Press.

Kiela, D., Hill, F., Korhonen, A., & Clark, S. (2014). Improving multi-modal representations using image dispersion: Why less is sometimes more. In *Proceedings of the 52nd annual meeting of the association for computational linguistics* (pp. 835–841).

Kousta, S.-T., Vigliocco, G., Vinson, D. P., Andrews, M., & Del Campo, E. (2011). The representation of abstract words: Why emotion matters. *Journal of Experimental Psychology: General*, *140*(1), 14–34.

Louwerse, M. M. (2011). Symbol interdependency in symbolic and embodied cognition. *Topics in Cognitive Science*, *3*, 273–302.

Louwerse, M. M., & Connell, L. (2011). A taste of words: Linguistic context and perceptual simulation predict the modality of words. *Cognitive Science*, *35*, 381–398.

Mandera, P., Keuleers, E., & Brysbaert, M. (2017). Explaining human performance in psycholinguistic tasks with models of semantic similarity based on prediction and counting: A review and empirical validation. *Journal of Memory and Language*, *92*, 57–78.

Marelli, M., & Baroni, M. (2015). Affixation in semantic space: Modeling morpheme meanings with compositional distributional semantics. *Psychological Review*, *122*(3), 485–515.

Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient estimation of word representations in vector space. In *Proceedings of workshop at the international conference on learning representation (ICLR).*

Mitchell, T. M., Shinkareva, S. V., Carlson, A., Chang, K.-M., Malave, V. L., Mason, R. A., & Just, M. A. (2008). Predicting human brain activity associated with the meanings of nouns. *Science*, *320*, 1191–1195.

Pennington, J., Socher, R., & Manning, C. D. (2015). GloVe: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)* (pp. 1532–1543).

Recchia, G., & Louwerse, M. M. (2015). Reproducing affective norms with lexical co-occurrence statistics: Predicting valence, arousal, and dominance. *The Quarterly Journal of Experimental Psychology*, *68*, 1584–1598.

Riordan, B., & Jones, M. N. (2011). Redundancy in perceptual and linguistic experience: Comparing feature-based and distributional models of semantic representation. *Topics in Cognitive Science*, *3*, 303–345.

Silberer, C., Ferrari, V., & Lapata, M. (2017). Visually grounded meaning representations. *IEEE Transactions on Pattern Recognition and Machine Intelligence*, *39*(11), 2284–2297.

Taboada, M. (2016). Sentiment analysis: An overview from linguistics. *Annual Review of Linguistics*, *2*, 325–347.

Tillmand, R., & Louwerse, M. (2018). Estimating emotions through language statistics and embodied cognition. *Journal of Psycholinguistic Research*, *47*(1), 159–167.

Utsumi, A. (2015). A complex network approach to distributional semantic models. *PLoS ONE*, *10*(8), e0136277.

Vecchi, E. M., Marelli, M., Zamparelli, R., & Baroni, M. (2017). Spicy adjectives and nominal donkeys: Capturing semantic deviance using compositionality in distributional spaces. *Cognitive Science*, *41*, 102–136.

Vigliocco, G., Kousta, S.-T., Rosa, P. A. D., Vinson, D. P., Tettamanti, M., Devlin, J. T., & Cappa, S. F. (2014). The neural representation of abstract words: The role of emotion. *Cerebral Cortex*, *24*, 1767–1777.

Zwaan, R., & Yaxley, R. (2003). Spatial iconicity affects semantic relatedness judgments. *Psychonomic Bulletin & Review*, *10*(4), 954–958.