

Yair Neuman*, Newton Howard, Louis Falissard and Rafi Malach

The embodiment of connotations: A proposed model

DOI 10.1515/sem-2016-0112

Abstract: The idea that abstract words are grounded in our sensorimotor experience is gaining support and popularity, as observed in the increasing number of studies dealing with “neurosemantics.” Therefore, it is important to form models that explain how to bridge the gap between basic bodily experiences and abstract language. This paper focuses on the embodiment of connotations, such as “sweet” in “sweet baby,” where the adjective has been abstracted from its concrete and embodied sense. We summarize several findings from recent studies in neuroscience and the cognitive sciences suggesting that *emotion*, *body*, and *language* are three factors required for understanding the emergence of abstract words, and (1) propose a model explaining how these factors contribute to the emergence of connotations, (2) formulate a computational model instantiating our theoretical model, and (3) test our model in a task involving the automatic identification of connotations. The results support our model pointing to the role of embodiment in the formation of connotations.

Keywords: embodiment, neurosemantics, connotations, computational semiotics

1 Introduction

Human beings have a unique ability to think “abstractly” in a way which seems to be detached from the realm of concrete sensorimotor experience. For example, the meaning of abstract words such as God, justice, and idea, cannot be trivially reduced to concrete perceptions or operations such as smelling or grasping.

Nevertheless, a current embodied approach in neuroscience and cognitive sciences rejects the idea that abstract concepts and words pop-up in our mind

***Corresponding author: Yair Neuman**, Department of Education, Ben-Gurion University of the Negev, Beer-Sheva, Israel, E-mail: yneuman@bgu.ac.il

Newton Howard: E-mail: newton.howard@nds.ox.ac.uk, **Louis Falissard:**

E-mail: louis.falissard@gmail.com, Nufield Department of Surgical Sciences, University of Oxford, Oxford, UK

Rafi Malach, Neurobiology, Department of Neurobiology, Weizmann Institute of Science, Rehovot, Israel, E-mail: rafi.malach@gmail.com

ex nihilo, and strives to ground (rather than to reduce) abstraction in sensorimotor experience (e.g., Borghi and Binkofski 2014; Harnad 1990; Lakoff 2014; Pulvermüller 2013; Shebani and Pulvermüller 2013). Although the exact meaning of embodiment is far from trivial (Kyselo and Di Paolo 2013), here we adopt the general meaning of the embodied approach to “ground” higher cognitive functions in sensorimotor activity.

1.1 Language and embodiment

It is possible to approach natural language as involving the processing of categories that are the essence of cognition (Harnad 2005). For instance, when a baby learns to point to his pet and to name it as a “Cat,” s(he) approaches it as a member of an abstract and general category, which is the Cat category. While the noun “cat” is considered to be a lexical category, we should recall that it is basically a semantic category describing a concrete object (Moseley and Pulvermüller 2014). The Cat category involves a different *logical type* than the one of the unique, individual and concrete cats from which it has been constructed. However, as signs in natural language are the names of categories, and categories are abstractions derived from concrete sensorimotor experiences with unique and individual entities (i.e., we never interact with an abstract cat), it is reasonable to conclude that natural language, despite its abstract semantic dimension, is grounded through categorization in concrete sensorimotor experience.

It is suggested that cognition “is fundamentally action bound, subserving the planning, selection, anticipation, and performance of actions” (Engel et al. 2013: 206). In this context, the “*embodiment of semantics*” describes theoretical attempts to ground the *meaning* of words in our sensorimotor experience. One such notable attempt is Lakoff and Johnson’s (1999) theory in which they propose *conceptual metaphors* to be the crucial mechanism for bridging the gap between concrete experiences and abstract concepts. A conceptual metaphor involves a mapping of relations from one conceptual domain to another conceptual domain. For example, when I would like to express my understanding of a certain issue, I may use the conceptual metaphor of understanding-as-seeing. In this case when I say, “I can see it,” I am using the domain of seeing as a source for describing my understanding. When I say, “I can see it,” I don’t use the verb “see” in its literal sense of perceiving with my eyes, but in an extended sense.

Others, such as Meteyard et al. (2012) suggest an integrative model according to which there are two main sources of information involved in the formation

of word meaning. These sources are described as “*experiential*” (i.e., sensory, motor, and affective) and “*linguistic*” (i.e., co-occurrence and syntactic patterns of words). For example, while the basic and simple meaning of “sweet” is grounded in our taste of food rich in carbohydrates, abstract words such as “justice” gain their meaning from semantic networks of words that appear with them (i.e., collocated with them) in our mental lexicon (Danesi 2003). For illustrating the linguistic dimension, we may examine the words that are collocated with an abstract word such as “freedom” in one of the most comprehensive linguistic corpora (Corpus of Contemporary American English [COCA]; Davies 2009). We find that the words collocated with freedom are words such as “democracy,” “expression,” and so on. Therefore, when asked what is the meaning of “freedom,” we may evoke a network of words and phrases, such as the “freedom of expression,” that load it with meaning.

The need to differentiate between the experiential and the linguistic factors of abstract words is grounded in the argument that sensorimotor simulation in itself cannot handle all language processing (Hagoort and Indefrey 2014). In this paper, we adopt this position but present a model showing how the linguistic and the experiential dimensions are related in the formation of connotations. This integrative conception of the experiential and linguistic dimensions is not new. In fact, it was shown by Vincent-Lamarre et al. (2016) that the meaning of abstract words in the lexicon could be reduced to a minimal set of words that are clearly concrete and embodied. In other words, while abstract words seem to gain their meaning only from a close network of associated signs, these signs in themselves draw on a minimal set of the network, which is clearly “experiential, i.e., grounded in our sensorimotor experience. For example, if we examine the definition of God, we can see that it includes words such as “creator,” “ruler,” and “universe.” These are abstract words, but we can recursively examine the definition of each word and find whether it has an embodied anchor. When following this procedure, we find that the definition of “creator” includes “person” which in its turn includes phrases such as “human being” that is defined in its turn by using more concrete words such as “man,” “woman” or “child.” In sum, even abstract words such as God are defined through words that are more concrete or defined through more concrete and embodied words. However, semantics involve a quantum leap of meaning when we move from the basic embodied set of words with their concrete references to more abstract words. In this context, it is important to better understand how this leap is taken.

1.2 The three threads of abstract words

Integrating the findings of the literature dealing with embodied semantics, we may find that it provides three different, albeit converging, threads associated with the meaning of abstract words.

The first thread involves the *affective dimension of abstract words*. Although it is argued that concrete words are processed more quickly, this finding is highly contextualized (Xiao et al. 2012). Kousta et al. (2011) present evidence to explain the speed of processing abstract words in terms of *emotional valence*. It is argued, that the faster cognitive processing of abstract words is due to the fact that their affective value is higher than that of concrete words. This thesis is supported by Vigliocco et al. (2014) who found that the processing of abstract words involves a greater activation in the rostral anterior cingulate cortex.

These studies, in addition to others, raise an interesting question: how is it that abstract words seem to be more emotionally loaded than concrete words that are more directly related to our sensorimotor experience? After all, it could have been expected that the emotional valence of a word or a phrase will be a simple function of its corresponding signified level of concreteness, such as in the case where the high negative valence of the word “snake” is a function of the concrete and potentially dangerous reptile it signifies. This is true in the case of highly emotional concrete objects but not necessarily true on average for such objects (e.g., shoes, stone, table).

One possible answer to the above question is that abstraction is a mechanism for *learning generalization*, and that it is inevitable that experiences loaded with stronger emotional valence will be generalized and abstract. In sum, both concrete and abstract words can be emotionally loaded but the finding that abstract words are generally more loaded may be explained through their grounding in and evolution from emotionally-loaded experiences they aim to generalize.

The second thread of abstract words involves the role of the body in the formation of meaning in the sense that different body parts as well as different sensory modalities are associated with concrete words and their extension to more abstract concepts. For example, Maouene et al. (2011) found that verbs that are more highly associated with a single region of the body (e.g., eat) are among the first acquired in vocabulary development and that the associations with a specific region of the body may be a better predictor of whether a verb is learned earlier than its concreteness or abstractness. Therefore, we should expect that in adjective-noun pairs (e.g., sweet baby) that involve the connotation of an adjective (e.g., “sweet” in the sense of “pleasant”) the embodied sense of the modified noun will have a lesser association with a specific sense than a noun

modified in a non-metaphorical expression, for instance “sweet ice-cream” where “ice-cream” is clearly associated with taste and tongue/mouth. This hypothesis will be tested in our experiment.

The third thread involved in the emergence of abstract words, concerns the *role of language* in the formation of abstract meaning (e.g., Danesi 2003; Dove 2014; Hoffman 2015). Wang et al. (2010) argue that abstract concepts produce greater activity in the inferior frontal gyrus (associated with language comprehension) and middle temporal gyrus (associated with word meaning) compared to concrete concepts, while concrete concepts produce greater activity in the posterior cingulate (associated with spatial memory), precuneus (associated with mental imagery), fusiform gyrus (associated with color information and face and body recognition), and parahippocampal gyrus (associated with scene recognition) compared to abstract concepts. These results suggest greater engagement of the *verbal system* for processing abstract concepts and therefore the emergence of abstract words seem to involve some kind of a quantum leap from a more basic perceptual experience to a more distributed network of words. In sum, networks of language and inferential processes, dissociated from a concrete experience seem to play a crucial role in the formation of abstract language (see Danesi 2003).

Neuman et al. (2012) also discuss the role of language in abstraction by drawing on Peirce’s idea of *Hypostatic Abstraction*. Peirce presented a mechanism according to which an experience, linguistically described in terms of an adjective (e.g., sweet) turns into an abstract object of contemplation expressed as a noun (e.g., sweetness). Neuman et al. (2012) provide empirical support for the existence of this mechanism and to the mediating role of language in the formation of this process.

What is the unique role of language in the formation of abstract meaning? The answer we propose is that by naming a memory of an embodied experience, we can extend the meaning of an experience and represent it as an autonomous representation, which is grounded-in albeit not reducible-to a concrete experience. For instance, when we learn the word “sweet” and associate it with a concrete attribute (i.e., TASTE), we can “graft” it (Danesi 2003) to other contexts that evoke a similar emotion, *which lack the original somatosensory experience*. In other words, we can develop abstract *connotations* (i.e., extended meanings) of the original experience through the use of words originally denoting a sensorimotor experience. This idea is developed in the proposed model.

2 The proposed model

Is there a way to explain how the body/senses, emotion, and language work together to generate the meaning of abstract words? Here we propose a simple

model, which is limited in its scope to the emergence of connotations involving adjective-noun phrases such as “dark thoughts,” “bitter person” and “deep mind.” Although our model is both developmental and evolutionary, the experiment that aims to support it is (1) located in the field of computational semantics and (2) limited for examining the role of the models’ factors in identifying connotations.

The model is presented and illustrated through a worked-out example concerning the connotation of “sweet.” The word “sweet” is an adjective clearly grounded in our sense of taste. When we taste sweet food the experience is associated with a certain sensory modality and body parts. The sweet taste also raises a strong positive emotion, which means that “sweet” has a positive *valence* and a high *arousal*. To recall, In Old English the meaning of “*swete*” was pleasing the senses. Now, when we encounter a new experience that activates the same valence and arousal, such as in the case of a mother observing her smiling baby, there is a potential for the activation of memories associated with of similar sense experiences. Therefore, the smiling baby who may evoke a high positive emotion may also evoke the embodied memory of sweetness. In other words, encountering a new experience that is similar in valence and arousal (i.e., affect) causes our memory (through sensory and cross-sensory activation) to enact the basic embodied experience of sweet taste not in a direct sense, but through memory and association only. However, and this is a highly important point, the sensory modalities associated with the baby in the example above are distinctly different from the modalities associated with the original experience of a sweet taste. That is, the positive and strong emotional experience evoked by the baby and the accompanied sensory and cross-sensory activation through memory is still not the original embodied experience of a sweet *taste*. According to this suggestion, a “sweet baby” is a baby evoking similar valence and arousal but at the same time lacking the original attributes of the sweet taste (i.e., the baby isn’t really sweet ...). Here we see how affect (positive/strong), body (sensory and cross-sensory activation), and language (the meaning of sweet is extended through networks of words i.e., sweet is “pleasing”) work together in producing the connotation. Like the Borromean Rings, you cannot remove one of these three processes without breaking apart the whole construct.

The formation of abstract connotations also involves a *violation of a habit*, a violation of “selectional preference” as proposed by Neuman et al. (2012). This hypothesized model urges us to think about abstract words as a product of an inferential process as well. This process is actually the abductive form of reasoning proposed by C. S. Peirce. Let us explain this point. When the human mind encounters a new experience evoking the arousal and valence of a sweet taste, it

may abductively infer that the experience is *actually* the experience of a sweet taste. Abduction is the process of “choosing a hypothesis” (CP 7.219) usually interpreted as the selection of the appropriate hypothesis given certain data, like a Bayesian form of reasoning. In our case, the data is the positive affect and the associated sensory experience. The hypothesis is the baby has a “sweet taste.” However, the generation of an abstract connotation *violates* this inferential process. If something new is hypothesized to be sweet but it *isn’t really sweet* in terms of sensorimotor experience then it may be either (1) not-sweet or (2) a different sense of sweet, an abstract form of sweetness. Take the common expression: “if it looks like a duck, swims like a duck, and quacks like a duck, then it is probably a duck.” We can think of this newly experienced sense of sweet as a “duck.” This “duck” provokes a similar affect to the original duck (i.e., the sweet taste) and along with it the cross-sensory activation that accompanied the experience of that first “duck” or sweet taste. However, regardless of these similarities what the mind experiences isn’t the duck, but metaphorically speaking a “shadow of the duck.” In the example of a “sweet baby,” the mind evokes the positive valence associated with a baby and associated with sweet. Then we ask whether a BABY is like the other concrete and embodied objects modified by sweet, mostly objects of the food category (e.g., cake) that we eat and taste. If the answer is negative (i.e., a baby is not an object in the food category) then the default decision is that it is sweet in its original embodied sense and extended sense (i.e., connotation) is formed.

The violation of basic embodied expectations as a mechanism of forming connotations is deeply related to the idea of *Selectional Preference* (Light and Greiff 2002), which describes the statistical “tendency” of words to co-occur with “selected” or “preferred” words from certain semantic categories. For example, when you read the word “sweet” your mind automatically expects that it will be followed by a word belonging to the *food* category. This powerful cognitive mechanism has been successfully used in the design of cognitively inspired intelligent systems for identifying abstract connotations (Neuman et al. 2013). Along the same line, when the experience of “sweet” is evoked in the context of a smiling baby, it involves *a violation of the expected semantic-embodied category*, because a baby is not food that we can taste. The phrase “a sweet baby” is therefore a violation of the selectional preference principle, albeit a violation that is grounded in a similar emotion and embodied experience. From this theorization we may hypothesize that in the case of adjective-noun phrases where the adjective is a connotation (e.g., dark humor), the modified noun (e.g., humor) will be less associated with the strongly embodied nouns that usually modify by the adjective (e.g., night, skin, hair, etc.). This is one of the hypotheses that will be tested in this paper.

2.1 Hypothesis 1

As previously proposed, connotations emerge when an object (e.g., a smiling baby) first raises a strong emotional valence similar to the one of the original denotation (e.g., sweet taste). If this hypothesis is grounded in adjective-noun phrases where the adjective functions in an abstract extended sense, we should find that the noun scores a higher valence than the noun modified by the adjective in “denotational” non-metaphorical phrases (e.g., dark hair).

2.2 Hypothesis 2

The second hypothesis concerns the role of the senses in the formation of connotations. We hypothesized that in “connotational” phrases comprised of adjective-noun pairs (e.g., sweet baby) the noun will (1) be less associated with a specific sense modality (e.g., taste) than the nouns in non-connotational phrases, and that; (2) there will be greater distance between the sense modalities of the adjective and the sense modalities of the noun. For instance, in the phrase “sweet cake” the noun cake is clearly associated with “taste,” which is the same sense modality associated with “sweet.” In contrast, in the context of “sweet dreams” the noun “dreams” will be more associated with “sight,” which is a different sense modality than the one associated with “sweet.”

2.3 Hypothesis 3

According to our model, when the mind has to decide whether “sweet” in “sweet dream” is a connotation, it first retrieves the concrete-embodied objects associated with (or in the case of language, modified by) “sweet” such as potato, chocolate, etc. Next the mind examines whether the modified noun belongs to the category/categories of the above-mentioned objects. For instance, in the context of “sweet dreams,” sweet usually modifies nouns from the FOOD category. Dream is not a FOOD. Therefore the conclusion is that “sweet dream” is a connotational phrase (i.e., the dream is not really sweet like sweet chocolate). Therefore, we have hypothesized that nouns in connotational phrases will be less associated with the categories of concrete/embodied objects usually modified by the target adjective.

3 Materials and method

If all of the above-mentioned factors contribute to the emergence of abstract language, we should identify their contribution in a computational model that aims to classify adjective-noun phrases as connotational or denotational.

3.1 The dataset

For testing the model, we have used the dataset prepared by Turney et al. (2011) for classifying metaphors as opposed to literal phrases, as connotations are tagged as Type III metaphors (Krishnakumaran and Zhu 2007).

The dataset is comprised of five adjectives with a clear embodied sense: *dark*, *deep*, *hard*, *sweet*, and *warm*. For each of the five adjectives, there are twenty word pairs in which the first word is an adjective and the second word is a noun. Each of the 100 phrases was accompanied by a rating by four judges whether it is metaphorical or literal. In 47 % of the phrases the adjective has been conceived by the judges as metaphorical and in the rest as literal (non-metaphorical). Although, Turney's et al. (2011) dataset is very limited in scope, we chose to use it as it has been intensively cited in the computational semantics literature as a kind of a gold-standard dataset for the task of automatically identifying adjective-noun metaphors.

4 Analysis and results

4.1 Testing hypothesis 1

Hypothesis 1 proposes that we should find that the noun scores higher on valence in connotational phrases. To test the first hypothesis, we adopted the distributional semantics approach for the computation of word meaning (Turney and Pantel 2010) and measured the valence of a noun by computing its distance from a vector of positively and negatively associated words. We used the norms of valence and arousal (Warriner et al. 2013), and selected the 30 highest/lowest ranked positive words. For instance, the positively loaded words are:

vacation happiness, happy, Christmas, enjoyment, fun, fantastic, lovable, free, hug, magical, delight, joy, joyful, relaxing, honest, sunshine, excited, accomplishment, bonus, comedy, courageous, laughter, lover, orgasm, cheerful, entertaining, love, oasis, peaceful

For measuring the distance between the vector of positively/negatively loaded words and each noun, we have used the Glove Suite (Pennington et al. 2014) and produced a measure of valence – Val – which is the average of the noun's similarity scores to the positive and negative vectors.

For testing the hypothesis, we have compared the Val score across the connotational and denotational phrases. As hypothesized, the nouns in the connotational phrases scored higher (0.13 versus 0.09, respectively). By using one-way ANOVA this difference was found to be statistically significant, $F(1, 98) = 8.92$, $p = 0.004$. However, we believe that a much more convincing evidence for supporting the model is to test it in a classification task. Therefore, we used the valence score as a predictor in a Binary Logistic Regression model with the phrase tag (connotational versus denotational) as a binary criterion, and applied a bootstrapping procedure of 1,000 samples. The model successfully predicted the phrases' tag, $\chi^2 = 8.73$, $p = 0.003$, with 67 % Accuracy, 55 % Recall, and 68 % Precision. Precision is a measure indicating how many of the cases the model identified as being connotations are actual connotations. In this sense, precision is an indication of how much we have improved in predicting whether a phrase is connotational given the model's prediction, over the prediction provided by the base-level of connotational phrases in our dataset. The Binary Logistic Regression model produced 68 % precision, a significant improvement over the base rate of connotational phrases in our dataset (47 %).

4.2 Testing hypothesis 2

Hypothesis 2 suggests that nouns in connotational phrases (a) have a lesser association with a specific sense than the nouns modified in denotational phrases and that in the case of connotational phrases (b), there will be greater distance between the sense modalities of the adjective and the sense modalities of the noun.

To test this hypothesis, we have used the Sensicon Lexicon¹ (Tekiroğlu et al. 2014) a resource associating English words with each of the five senses: Sight, Hearing, Taste, Smell and Touch. To test hypothesis 2a, we have identified the sense that *was judged to be the most associated* with the target noun and used its score (MaxSense) for comparing nouns in the connotational and denotational phrases. As hypothesized, nouns in metaphorical phrases scored lower on the MaxSense measure (0.05 versus 0.10, respectively). Using one-way ANOVA, the difference was found to be statistically significant, $F(1, 98) = 25$, $p < 0.001$. In this

¹ We thank Serra Sinem Tekiroglu for providing us with the lexicon.

case, the Binary Logistic Regression was statistically significant, $\chi^2=22.45$, $p<0.001$, with 69 % Accuracy, 77 % Recall, and 64 % Precision.

For testing hypothesis 2b, we have measured the score of the adjective on each sense modality minus the score of its modified noun on the same sense and summed the results across the five senses. This measure (SumSense) is indicative of the degree in which the adjective and the noun differ in their corresponding sense modalities. Comparing SumSense across metaphorical and literal phrases, we have found a statistically significant difference, $F(1, 98)=9.50$, $p=0.003$. The distance between the sensual modalities of the adjective and the noun was found to be higher in metaphorical phrases (0.16 versus 0.05, respectively). These results suggest that the distance between the senses associated with the adjective and the noun can be used for identifying the connotational sense of a phrase. In this case, the Binary Logistic Regression model was found to be statistically significant, $\chi^2=9.19$, $p=0.002$, with 63 % Accuracy, 49 % Recall, and 64 % Precision.

4.3 Testing hypothesis 3

For testing the third hypothesis, we have identified the first 100 nouns modified by each of our adjectives in COCA. This procedure has been applied by automatically identifying in COCA the nouns collocated one place to the right of our target adjective. Next, we used Sensicon. To recall, each noun has five scores on each of the five senses. For each noun, we chose its highest score. Then we select the 30 nouns that scored highest, meaning we chose the 30 nouns with the highest embodied sense. Next, we generated a similarity matrix based on the vector measurement for each list of nouns associated with a given adjective. The descriptor “sweet” for instance, returned a matrix that included “potato” and “chocolate” and their semantic distance as measured through the distance between vectors. We applied Principle Component Analysis (PCA) and identified the three most significant factors underlying our data. These are the factors that organize the nouns in the most economical way. Finally, for each of the nouns in our adjective-noun phrases, we measured the distance between the noun and each of its relevant factors by computing the distance between the vector of the noun and the vectors of the words comprising the 3 factors identified by the PCA.

The scores describing the noun’s distance from each of the 3 major factors were titled cluster 1, cluster 2 and cluster 3, respectively. To test our hypothesis, we have used a one-way ANOVA in order to compare these scores across metaphorical and literal phrases. The difference was found to be statistically

significant only for cluster 1, $F(1,98) = 5.88$, $p = 0.017$, with connotational phrases showing a greater distance from the cluster (0.06 versus -0.06) as hypothesized. This result is trivial, as cluster 1 represents the most significant cluster of words modified by the adjective. The Regression model was statistically significant, $\chi^2 = 5.84$, $p = 0.02$, and produced 58 % Accuracy, 55 % Recall, and 55 % Precision. To measure the contribution of all the variables that were found to differentiate between connotational and denotational phrases, we have used the k-Nearest Neighbor classifier with ten-fold cross validation. In this case we have gained 73 % Accuracy, 61 % Recall, and 76 % Precision.

5 Discussion

Pulvermüller (2013) suggests that there are four neuro-semantic mechanisms: referential semantics, combinatorial semantics, emotional-affective semantics, and abstraction mechanisms. Referential semantics involve a simple correspondence between symbols and objects or actions. In this sense, we describe the meaning of the symbol/sign as “denotation.” Combinatorial semantics enables the learning of symbolic meaning from context. As the denotation, the first or the simple meaning of certain abstract terms involves contextual understanding, the mechanism of combinatorial semantics may be responsible for denotation in its disembodied sense. Emotional-affective semantics establishes links between signs and internal states of the body, and abstraction mechanisms generalize over a range of instances of semantic meaning. The model proposed in this study clearly corresponds with the mechanisms proposed by Pulvermüller, as our model suggests that when encountering an experience similar in valence/arousal to an embodied experience encountered in the past (e.g., sweet taste), embodied schemes are activated. This activation generates anticipation for a certain object/experience of the same embodied sense. This anticipation is actually the result of the habit grounded in a previous sensorimotor experience, such as when positive and strong affect evokes the expectation of taste and food. When this habit or anticipation is violated, language generates an abstract extension of the original experience through a differentiated abstract meaning (e.g., connotation) of the original sign (e.g., sweet) and its valence. This extension is a result of both a default mechanism (i.e., the realization that we don’t deal with the simple referential sense) and linguistic networks (Danesi 2003) that provide this default conclusion with a symbolic meaning.

While human language may be grounded in a more basic sensorimotor experience, it cannot be reduced to sensorimotor experience alone and so embodied

models of language and semantics should carefully avoid the reductionist pitfall and instead develop “emergentalist” models that respect the embodied nature of language as well as its complex nature. The experiment we have performed on a limited dataset of adjective-noun phrases cannot fully support the above-mentioned model. However, the experiment *provides preliminary evidence* to the role of the emotional, embodied, and linguistic in the identification, and hence in the formation, of connotations. These preliminary results invite extension into a much more comprehensive dataset. Second, while we point to the importance of the three-above-mentioned threads in the identification of connotations, this indirect support of our model has been established in the context of computational semantics and should be validated with human subjects and by carefully elucidating the neural mechanisms underlying the comprehension and generation of abstract connotations. Finally, the formation of abstract language, specifically the formation of abstract connotations, is both an evolutionary and a developmental process. This suggests that future research must pay close attention to the development of computational models simulating the emergence of human semantics at the evolutionary, cultural, and developmental levels of analysis. We therefore conclude by pointing to possible future directions of research and inviting researchers to examine these directions.

References

- Borghi, A. M. & F. Binkofski. 2014. *Words as social tools: An embodied view on abstract concepts*. New York: Springer.
- CP 7.219. 1933. *The collected papers of Charles S. Peirce*, 8 vols., C. Hartshorne, P. Weiss & A. W. Burks (eds.). Cambridge: Harvard University Press.
- Danesi, M. 2003. Metaphorical “networks” and verbal communication: A semiotic perspective on human discourse. *Sign Systems Studies* 31. 341–363.
- Davies, M. 2009. The 385+ million word Corpus of Contemporary American English (1990–2008+): Design, architecture, and linguistic insights. *International Journal of Corpus Linguistics* 14(2). 159–190.
- Dove, G. 2014. Thinking in words: Language as an embodied medium of thought. *Topics in Cognitive Science* 6(3). 371–389.
- Engel, A. K., A. Maye, M. Kurthen & P. König. 2013. Where’s the action? The pragmatic turn in cognitive science. *Trends in Cognitive Sciences* 17(5). 202–209.
- Hagoort, P. & P. Indefrey. 2014. The neurobiology of language beyond single words. *Annual Review of Neuroscience* 37. 347–362.
- Harnad, S. 1990. The symbol grounding problem. *Physica D: Nonlinear Phenomena* 42(1). 335–346.
- Harnad, S. 2005. To cognize is to categorize: Cognition is categorization. In C. Lefebvre & H. Cohen (eds.), *Handbook of categorization in cognitive science*, 20–45. Amsterdam: Elsevier.

- Hoffman, P. 2015. The meaning of “life” and other abstract words: Insights from neuropsychology. *Journal of Neuropsychology* 10(2). 317–343.
- Kousta, S. T., G. Vigliocco, D. P. Vinson, M. Andrews & E. Del Campo. 2011. The representation of abstract words: Why emotion matters. *Journal of Experimental Psychology: General*, 140(1). 14–34.
- Krishnakumaran, S. & X. Zhu. 2007. Hunting elusive metaphors using lexical resources. In *Proceedings of the Workshop on Computational approaches to Figurative Language*, 13–20. Association for Computational Linguistics.
- Kyselo, M. & E. Di Paolo. 2013. Locked-in syndrome: A challenge for embodied cognitive science. *Phenomenology and the Cognitive Sciences* 14. 517–542.
- Lakoff, G. 2014. Mapping the brain’s metaphor circuitry: Metaphorical thought in everyday life. *Frontiers in Human Neuroscience* 8. 958.
- Lakoff, G. & M. Johnson. 1999. *Philosophy in the flesh*. New York: Basic.
- Light, M. & W. Greiff. 2002. Statistical models for the induction and use of selectional preferences. *Cognitive Science* 26(3). 269–281.
- Maouene, J., N. Sethuraman, A. Laakso & M. Maouene. 2011. The body region correlates of concrete and abstract verbs in early child language. *Cognition, Brain, Behavior* 15(4). 339–383.
- Meteyard, L., S. R. Cuadrado, B. Bahrami & G. Vigliocco. 2012. Coming of age: A review of embodiment and the neuroscience of semantics. *Cortex* 48(7). 788–804.
- Moseley, R. L. & F. Pulvermüller. 2014. Nouns, verbs, objects, actions, and abstractions: Local fMRI activity indexes semantics, not lexical categories. *Brain and Language* 132. 28–42.
- Neuman, Y., D. Assaf, Y. Cohen, M. Last, S. Argamon, N. Howard & O. Frieder. 2013. Metaphor identification in large texts corpora. *PLoS One* 8(4). e62343.
- Neuman, Y., P. Turney & Y. Cohen. 2012. How language enables abstraction: A study in computational cultural psychology. *Integrative Psychological and Behavioral Science* 46(2). 129–145.
- Pennington, J., R. Socher & C. D. Manning. 2014. GloVe: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 1532–1543. Association for Cognitive Linguistics.
- Pulvermüller, F. 2013. How neurons make meaning: Brain mechanisms for embodied and abstract-symbolic semantics. *Trends in Cognitive Sciences* 17(9). 458–470.
- Shebani, Z. & F. Pulvermüller. 2013. Moving the hands and feet specifically impairs working memory for arm- and leg-related action words. *Cortex* 49(1). 222–231.
- Tekiroğlu, S. S., G. Özbal & C. Strapparava. 2014. Sensicon: An automatically constructed sensorial lexicon. In *Proceedings of the 2014 conference on Empirical Methods in Natural Language Processing (EMNLP)*, 1511–1521. Association for Cognitive Linguistics.
- Turney, P. D., Y. Neuman, D. Assaf & Y. Cohen. 2011. Literal and metaphorical sense identification through concrete and abstract context. In *Proceedings of the 2011 conference on the Empirical Methods in Natural Language Processing (EMNLP)*, 680–690. Association for Cognitive Linguistics.
- Turney, P. D. & P. Pantel. 2010. From frequency to meaning: Vector space models of semantics. *Journal of Artificial Intelligence Research* 37(1). 141–188.
- Vigliocco, G., S. T. Kousta, P. A. Della Rosa, D. P. Vinson, M. Tettamanti, J. T. Devlin & S. F. Cappa. 2014. The neural representation of abstract words: The role of emotion. *Cerebral Cortex* 24(7). 1767–1777.

- Vincent-Lamarre, P., A. B. Massé, M. Lopes, M. Lord, O. Marcotte & S. Harnad. 2016. The latent structure of dictionaries. *Topics in Cognitive Science* 8(3). 625–659.
- Wang, J., J. A. Conder, D. N. Blitzler & S. V. Shinkareva. 2010. Neural representation of abstract and concrete concepts: A meta-analysis of neuroimaging studies. *Human Brain Mapping* 31. 1459–1468.
- Warriner, A. B., V. Kuperman & M. Brysbaert. 2013. Norms of valence, arousal, and dominance for 13,915 English lemmas. *Behavior Research Methods* 45(4). 1191–1207.
- Xiao, X., D. Zhao, Q. Zhang & C. Y. Guo. 2012. Retrieval of concrete words involves more contextual information than abstract words: Multiple components for the concreteness effect. *Brain and Language* 120(3). 251–258.