Predicting the Good Guy and the Bad Guy:
Attitudes are Encoded in Language Statistics

Gabriel Recchia[1] (grecchia@memphis.edu)
Alexandra L. Slater[1] (alslater@memphis.edu)
Max M. Louwerse[1,2] (mlouwerse@uvt.nl)

[1]Institute for Intelligent Systems, University of Memphis
365 Innovation Drive, Suite 303, Memphis, TN 38152, USA

[2]Tilburg Center for Cognition and Communication, Tilburg University
Warandelaan 2, 5037 AB Tilburg, The Netherlands

Abstract

Various studies have provided evidence that people activate introspective simulations when making valence judgments. Such evidence is in line with an embodied cognition account that argues that cognition is fundamentally embodied, with perceptual simulation rather than language statistics being the source of lexical semantics. Recently, demonstrations that conceptual knowledge is encoded in language have been used to argue that semantic processing involves both language statistics and perceptual simulation, with linguistic cues allowing meaning to be bootstrapped with minimal symbol grounding. Whether language also encodes attitudes towards concepts is unclear. In three studies, negative-valence words were found to be more closely associated in language with individuals commonly considered villains, and positive-valence words with heroes (both fictional and historical). These results suggest that attitudes toward persons can be inferred from lexical associations.

Keywords: affective norms; emotion; valence; latent semantic analysis; embodied cognition; distributional semantics

Introduction

Embodied cognition accounts emphasize that language evokes perceptual simulations (Barsalou, 1999; Glenberg, 1997; Pecher & Zwaan, 2005; Semin & Smith, 2002). For instance, a word like *eagle* automatically activates the visual system, whereby we 'see' the eagle in the sky (Pecher, van Dantzig, Boot, Zanzolie, & Huber, 2010, Šetić & Domijan, 2007). Similarly, words like *lick, pick,* and *kick* automatically activate regions of the motor system associated with the tongue, hand, and foot, respectively (Pulvermüller, 2005). For words with affective content, these simulations are presumed to incorporate "introspective simulations" (Barsalou, 1999) or "affective images" (Paivio, 2013), re-enactments of emotional states. An alternative account argues that perceptual content may be encoded in language statistics as well, allowing comprehenders to use distributional semantics to retrieve perceptual information without always resorting to simulation (Louwerse, 2011).

Various studies have argued that the valence of words is perceptually simulated (Meier & Robinson, 2004; Meier, Hauser, Robinson, Friesen, & Schjeldahl, 2007; Pecher et al., 2010). For instance, when the word *joy* is presented on the top of the screen (and *hate* is presented on the bottom of the screen) it is processed faster and remembered better because of perceptual simulation (Meier & Robinson, 2004). Van Dantzig, Zeelenberg and Pecher (2009) asked participants to move valenced words (e.g. *despair* or *pleasure*) toward or away from another valenced word (e.g., *reward* or *revenge*) and demonstrated an embodied approach/avoidance effect. In short, embodied cognition studies have demonstrated that concrete words and abstract words, such as valence words, activate perceptual simulations.

In previous work we have demonstrated that perceptual information is encoded in language (Louwerse, 2008; Louwerse & Benesh, 2012), so that it is these linguistic associations rather than perceptual simulations per se that might trigger effects that have been attributed to embodied cognition (Hutchinson & Louwerse, in press; Louwerse & Connell, 2011; Louwerse & Jeuniaux, 2010). For instance, the fact that *eagle-dolphin* is processed faster than *dolphin-eagle* might be explained by "high" words typically preceding "low" words in language (cf., *up and down*, *top and bottom*, *head and shoulders*). If language statistics can explain processing concrete words (Louwerse & Jeuniaux, 2010), we would expect that it can also explain processing abstract words, such as valence words.

Indeed, evidence from dual coding theory (Paivio, 2010) suggests that the role of language may be particularly strong in the case of abstract words. Associations between abstract words and other linguistic symbols (e.g., words and grammatical constituents) may thus drive processing even more strongly than associations with introspective simulations, particularly early in processing, when representations of linguistic forms are most active (Barsalou, Santos, Simmons & Wilson, 2008; Louwerse & Jeuniaux, 2008). For example, Paivio (1978) found that pleasantness ratings were fastest in response to images, second-fastest in response to concrete words, and slowest in response to abstract words. Paivio hypothesized that this pattern is observed because concrete concepts are linked directly to embodied affective information, whereas abstract words are linked more strongly to other linguistic

information and only indirectly to embodied affective information.

Affective simulations of linguistic stimuli may therefore not always be constructed in tasks requiring affective judgments. This view is consistent with Tillman, Hutchinson, Jordan, & Louwerse (2013), who investigated switching costs produced by verifying affective properties – e.g., the increase in the amount of time required to process a happy sentence ("birthdays can be happy") when it was preceded by a sad sentence ("insults can be devastating"), in contrast to when it was preceded by another happy sentence. They found that linguistic association (first-order co-occurrences in text) was a better predictor of fast reaction times than affective priming (inducing subjects to smile or frown), but that the reverse was true for slow reaction times.

However, if comprehenders can gain information about the valence of a word from the valence of its linguistic associates, language needs to encode this information. Evidence in favor of this hypothesis comes from Bestgen & Vincze (2012), who were able to predict the valence of 1,034 words in the ANEW affective norms ($r = .71$) on the basis of the valence of words with which they were associated in text. However, Klauer & Musch (2001) note that not all words with similar affective properties are associated in language, and were unable to obtain affective priming effects between words that shared similar levels of valence but were not linguistically associated (e.g. *sunshine* and *loyalty*). A particularly intriguing question is whether liked/disliked persons are linguistically associated with positively/negatively valenced words.

Study 1 and 2 focused on predicting heroes and villains in fictional texts. In Study 1, we investigated whether valence of fictional characters could be predicted from linguistic associations in a set of novels in which they appeared, the Harry Potter series. Study 2 extended the findings from Study 1 to the question of whether the valence of fictional characters could be predicted from the text of Wikipedia. Finally, in Study 3 we investigated whether these findings could be extended to historical figures. Using Wikipedia text provided a particularly strong test of the hypothesis, as one of its founding principles is that articles be written in a "neutral point of view," that is, "representing fairly, proportionately, and, as far as possible, without bias, all of the significant views that have been published by reliable sources on a topic" (Wikipedia, 2013). Based on the work we have done showing that language encodes perceptual information (Louwerse & Jeuniaux, 2010), including valence information (Tillman, et al., 2013), we predicted that language statistics allows for attitudes toward persons to be estimated from lexical associations, and that disliked individuals co-occur with negative-valence words, and liked individuals co-occur with positive-valence words, even in texts deliberately written with a neutral point of view.

Study 1

In Study 1, we investigated whether valence of fictional characters could be predicted from linguistic associations. The Harry Potter series was chosen partly due to its unambiguous identification of groups of "good" and "evil" characters, establishing a clear ground truth for evaluation. Characters are easy to classify by their membership in one of four groups. The Order of the Phoenix and Dumbledore's Army are comprised of "good" characters with positive moral attributes, while Death Eaters and the Inquisitorial Squad are comprised of "evil" characters with negative moral attributes. We hypothesized a crossover interaction such that good characters would be more closely related to positive-valence words than to negative-valence words, whereas the reverse would be true for evil characters.

Following Bestgen & Vincze (2012), we used Latent Semantic Analysis (LSA; Landauer & Dumais, 1997) to quantify the degree of association between the words under investigation (in our case, character names) and words of known valence. LSA is commonly used in psychology, computational linguistics, and information retrieval to quantify the degree of linguistic association between words. Its estimates of similarity between word meanings have achieved scores on the synonymy section of the Test of English as a Foreign Language that rival human performance (Landauer & Dumais, 1997), and it has been successfully applied to tasks as diverse as assessing reading comprehension (Foltz, Kintsch, & Landauer, 1998) and simulating human word association norms (Steyvers, Shiffrin, & Nelson, 2004; Jones, Gruenenfelder, & Recchia, 2011). LSA takes as input a matrix for which the value (i, j) of each cell indicates the number of times word i occurs in document j. Each term is weighted so as to reduce the influence of very frequent words, and singular value decomposition is applied to factor the matrix into three new matrices U, S, and V' whose product yields the original matrix. By truncating to a fixed number of dimensions prior to computing the product, a new matrix of lower rank can be obtained. This serves as a low-dimensional approximation of the original matrix. Finally, the similarity between two words can be obtained by computing the cosine between their corresponding rows. LSA cosines therefore yield a text-based measure of second-order linguistic association.

Method

Lists of good and evil characters were obtained from two separate fan based encyclopedias: the Harry Potter Wiki (http://harrypotter.wikia.com) and The Harry Potter Lexicon (http://www.hp-lexicon.org). The Harry Potter Wiki is a fan-based community wiki, while the Harry Potter Lexicon is an online encyclopedia of topics related to the Harry Potter series. Characters' allegiance to either good groups (The Order of the Phoenix and Dumbledore's Army) or evil groups (Death Eaters and the Inquisitorial Squad) listed on each site were nearly identical, and all characters possessing a proper name were included. The good character list

consisted of 61 unique characters, while the evil list included 48.

Words with positive and negative valence were obtained from the original version of the ANEW norms (Bradley & Lang, 1999) with 1,034 words. A median split was performed on the valence values. All words above the median (5.295) were labeled high-valence words, while those below the median were labeled low-valence. An LSA space of 300 dimensions was constructed for the text of the Harry Potter series according to the method recommended by Quesada (2006). Paragraphs were used as documents, and cosines were computed between each character name and each word in the ANEW norms. Finally, LSA cosines between all pairings of character names and words were subjected to a two-way analysis of variance having two levels of integrity (good, evil) and two levels of valence (high-valence words, low-valence words).

Results and Disccusion

Mixed effects models were run on the LSA cosine values. Because of the nature of mixed effects models and the large degree of freedom, F-test denominator degrees of freedom were estimated using the Satterthwaite degrees of freedom adjustment to reduce the chance of Type I error.

There was no main effect of integrity, $F(1, 112702) = .207$, $p = .65$, nor was there a main effect of valence, $F(1, 112702) = .002$, $p = .96$. Importantly, however, there was a significant interaction (Figure 1), $F(1, 112702) = 437.5$, $p < .001$, one-tailed (Figure 1). Specifically, names of evil characters had higher cosines to low-valence words ($M = .0035$, $SD = .0516$) than to high valence words ($M = .0026$, $SD = .0516$), $t(49630) = -1.89$, $p = .03$, one-tailed. Names of good characters showed the reverse pattern, having higher cosines to high-valence words ($M = .0034$, $SD = .0457$) than to low-valence words ($M = .0025$, $SD = .0449$), $t(63072) = 2.36$, $p = .01$, one-tailed.

We observed a robust interaction consistent with the hypothesis. However, this is not particularly surprising, given that characters in Harry Potter are described with morally loaded adjectives and verbs to convey their moral alignment. In Study 2, we therefore investigated whether similar results could be obtained using an alternative set of characters and a more encyclopedic text.
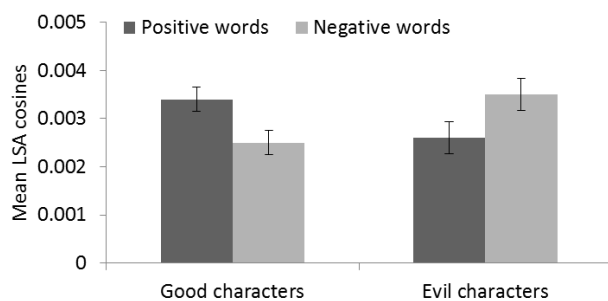

Figure 1. Crossover interaction between integrity and valence (Study 1).

Study 2

Study 2 aimed to validate the results of Study 1 by examining whether similar results could be obtained using an alternative set of fictional characters on a less explicitly biased text. Furthermore, Study 1 was conducted only using one set of affective norms, namely, the original ANEW list (Bradley & Lang, 1999). These norms have since been expanded to a list containing 2,471 unique words (Bradley & Lang, 2010), and a far more extensive set of affective norms has been independently collected by Warriner, Kuperman, and Brysbaert (2013). If the effect found in Study 1 is robust, the same effect should be observed irrespective of the particular set of norms used. We again hypothesized a crossover interaction analogous to that observed in Study 1 for both sets of norms.

Method

A list of 100 iconic heroes and villains in American cinema was obtained from the American Film Institute's *100 Years… 100 Heroes and Villains*, a list of 100 movie characters rated by expert judges as having left a particularly distinctive cultural impact and cinematic legacy. A hero was defined by the Institute as a character "who prevails in extreme circumstances and dramatizes a sense of morality, courage and purpose," while villains were defined as "ultimately tragic" characters exhibiting wickedness of mind and selfishness of character (American Film Institute, 2003). Constructing an LSA space for of the entire Wikipedia corpus proved infeasible due to computational limitations, so the LSA space was constructed using the subset of Wikipedia consisting of all documents that contained the name of any hero or villain on the list. As in Study 1, character names were treated as single tokens in the LSA space, and character pairs (e.g., Thelma Dickerson & Louise Sawyer) were treated as distinct characters. One character appeared on both lists (the Terminator, a hero in *Terminator* and a villain in *Terminator 2*) and was omitted from the analysis. Low- and high-valence words were computed separately for each set of norms on the basis of a median split, with single-token words above the median constituting the set of high-valence words, and those below it the low-valence words. ANOVAs were conducted using both the ANEW norms and the affective norms of Warriner et al. as in Study 1. Because over 95% of the instances in Wikipedia for one character (*Man*, the villain in Bambi) did not refer to the villain in question, ANOVAs were first computed with *Man* excluded from the set of villains, and then again with *Man* included.

Results and Discussion

The statistical analysis was identical to the one performed in Study 1, in terms of independent and dependent variables and the Satterthwaite degrees of freedom adjustment. For the analysis conducted on the ANEW norms, there was a main effect of perceived integrity, $F(1, 199835) = 61.2$, $p < $

.001, with words having higher cosines to names of villains (M = .0023, SD = .0601) than to names of heroes (M = .0002, SD = .0572). There was also a main effect of valence, $F(1, 199835)$ = 11.3, p = .001, such that character names had higher cosines to low-valence words (M = .0016, SD = .0586) than to high-valence words (M = .0008, SD = .0587). The interaction effect (Figure 2) was also significant, $F(1, 199835)$ = 4.7, p = .015 (one-tailed). Specifically, names of villains had higher cosines to low-valence words (M = .0030, SD = .0603) than to high-valence words (M = .0016, SD = .0599), $t(94175)$ = 3.7, p < .001, but names of heroes were not any more similar to low-valence words than to high-valence words, $t(105660)$ = .89, p = .37, likely due to the very low LSA cosine values (Figure 2).

Similar results were obtained for the analysis conducted on the Warriner et al. norms. There was a main effect of perceived integrity, $F(1, 1022072)$ = 65.6, p < .001, with words having higher cosines to names of villains (M = .0015, SD = .0603) than to names of heroes (M = .0006, SD = .0582). There was also a main effect of valence, $F(1, 1022072)$ = 62.8, p < .001, such that character names had higher cosines to low-valence words (M = .0015, SD = .0599) than to high-valence words (M = .0006, SD = .0586). Finally, the interaction effect was significant, $F(1, 1022072)$ = 12.0, p < .001 (one-tailed). As before, villains had higher cosines to low-valence words (M = .0022, SD = .0611) than to high-valence words (M = .0008, SD = .0595), $t(481666)$ = 7.7, p < .001. Heroes also had higher cosines to low-valence words (M = .0008, SD = .0588) than to high-valence words (M = .0003, SD = .0577), $t(540406)$ = 3.3, p < .001, but to a far lesser extent than was true for villains, with the mean difference in cosine similarities between high- and low-valence words being 2.5 times greater for villains than for heroes. Repeating the ANEW and Warriner et al. ANOVAs with "Man" included as a villain did not change the directionality or presence of any effect[1].

On the basis of the results from Study 1, we hypothesized a crossover interaction such that heroes would be more closely related to positive-valence words than to negative-valence words, whereas the reverse would be true for villains. Although villains were indeed more similar to negative than positive words, no difference was observed for heroes. In addition, it was unclear why main effects were observed in Study 2 and not in Study 1. It is unlikely that these are due to differences in frequency, as the explicit purpose of the log-entropy weighting step of LSA is to correct for differences in term count. Furthermore, although both main effects suggested that both sets of negative items (low-valence words and names of villains) are more similar
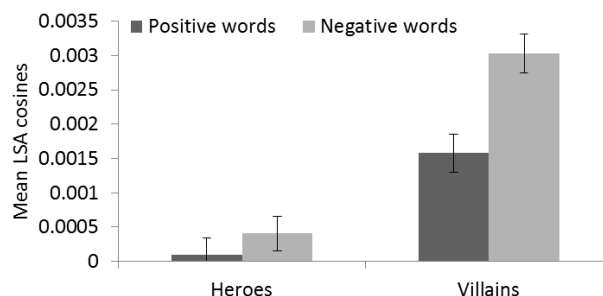

Figure 2. Interaction between perceived integrity and valence, ANEW norms (Study 2).

to terms in general than are positive items (high-valence words and names of heroes), high-valence words tend to be more frequent in English than low-valence words in both the ANEW and Warriner et al. norms, whereas names of heroes were if anything less frequent than names of villains, though not significantly so (p = .7). In any event, they are not directly relevant to the hypothesis under investigation.

Studies 1 and 2 considered fictional characters. This was useful for establishing that unambiguously negative individuals are more strongly associated with negative words (in fiction and in encyclopedic text), but does not tell us whether such effects are likely to apply to actual people. In Study 3, we aimed to replicate Study 2 with sets of positively and negatively perceived historical figures rather than fictional characters.

Study 3

The goal of Study 3 was to determine whether the results of Studies 1 and 2 remained valid for real individuals rather than fictional characters. If so, the results of Tillman et al. (2013) discussed earlier would imply that linguistic associations may plausibly be relied upon as a short-cut when we evaluate the valence of individuals. Based on the results of Studies 1 and 2, we predicted that historical villains would be more closely associated with low-valence than high-valence words but that the reverse would be true for historical heroes (as in Study 1), as well as a weaker hypothesis – based on the finding of Study 2 – that only historical villains would be more closely associated with low-valence words.

Method

Study 3 followed the same protocol as Study 2, with the list of movie heroes replaced with the eighteen individuals identified by the Gallup Organization as *Gallup's List of People that Americans Most Widely Admired in the 20th Century* (Gallup, 1999). The list of villains was replaced with the eighteen individuals topping the list of *The All-Time Worst People in History*, a list created by a continuous online poll on which over 14,000 individuals had voted at the time of retrieval (Ranker, 2013). All other aspects of the methodology were conducted as described in Study 2.

---

[1] Repeating the ANOVAs with *Man* included as a villain yielded same-direction main effects of perceived integrity and valence, all at p < .001. Analogous interaction effects were also obtained (ANEW, p = .02, one-tailed; Warriner et al., p < .001, one-tailed), with mean difference in cosine similarities between high- and low-valence words greater for villains than for heroes (ANEW, 4.4 times greater; Warriner et al., 2.4 times greater).

Results

The statistical analysis and its parameters were identical to Study 1 and Study 2. Using the ANEW norms, there was once again a main effect of perceived integrity, $F(1, 73688) = 6.37$, $p = .01$, with words having higher cosines to names of villains ($M = .0026$, $SD = .0557$) than to names of heroes ($M = .0016$, $SD = .0540$). There was no main effect of valence, $F(1, 73688) = .90$, $p = .35$. As in the analysis on movie characters, there was a significant interaction, $F(1, 73688) = 18.38$, $p < .001$ (one-tailed), such that villains had higher cosines to low-valence words ($M = .0033$, $SD = .0581$) than to high valence words ($M = .0019$, $SD = .0535$), $t(36844) = 2.3$, $p = .02$ (Figure 3). Names of heroes showed the reverse pattern, having higher cosines to high-valence words ($M = .0026$, $SD = .0549$) than to low-valence words ($M = .0005$, $SD = .0531$), $t(36844) = -3.8$, $p < .001$. There were no significant main effects in the ANOVA conducted using the Warriner et al. norms, but there was a significant interaction, $F(1, 346820) = 66.6$, $p < .001$ (one-tailed), such that villains had higher cosines to low-valence words ($M = .0023$, $SD = .0562$) than to high valence words ($M = .0010$, $SD = .0523$), $t(173410) = 5.1$, $p < .001$. Names of heroes showed the reverse pattern, having higher cosines to high-valence words ($M = .0027$, $SD = .0548$) than to low-valence words ($M = .0010$, $SD = .0546$), $t(173410) = -6.4$, $p < .001$.
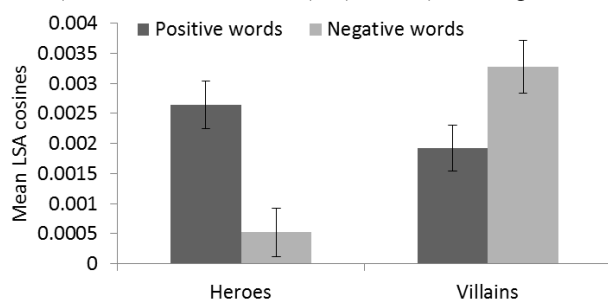
Figure 3. Interaction between perceived integrity and valence, ANEW norms (Study 3).

General Discussion

Given evidence that linguistic units associated with concrete and abstract terms may play a part in evaluative processes, it is of interest whether the information required to form evaluative judgments of persons is encoded in linguistic associations. In three different studies conducted over multiple sets of affective norms, negative-valence words were found to be more closely associated in language with individuals commonly considered villains and positive-valence words were more closely associated with individuals embodying heroic attributes (at least for historical figures and fictional characters). These results suggest that attitudes toward persons can be inferred from lexical associations, even from texts deliberately written with a "neutral point of view." In other words, sufficient information is available from linguistic statistics to make valence judgments without resorting to perceptual or affective simulation. The degree to which linguistic associations can or cannot account for accounts of embodied effects in valence processing has yet to be investigated, but these results suggest that the potential impact of linguistic associations should not be ignored. It is also notable that this paper found significant effects even when using a very simple, co-occurrence-based measure of association. It is well-known that LSA is not sensitive to elements of meaning that require attention to linguistic structure, such as negation, anaphora, and semantic roles. More sophisticated algorithms may be able to make even better use of the statistics encoded in language. Furthermore, computing the valence of an individual's linguistic associates may be of use to social scientists, e.g., as a method for estimating the degree to which particular public figures are described in positive or negative terms and how this has changed over time, or obtaining a quantitative evaluation of the degree of bias with which a text speaks about a particular individual.

Whether comprehenders utilize linguistic associations in the formation of their attitudes has not been investigated within the scope of this paper. Based on our other research, our prediction is that depending on the cognitive task (Louwerse & Jeuniaux, 2010), the time course (Louwerse & Connell, 2011; Louwerse & Hutchinson, 2012), and individual differences (Hutchinson & Louwerse, in press), comprehenders use these statistical linguistic cues in their comprehension processes, for which the reported encoding of affective information in language statistics is a prerequisite.

References

American Film Institute. (2003). *AFI's 100 Years... 100 Heroes and Villains*. Retrieved from AFI.com: http://www.afi.com/100years/handv.aspx

Barsalou, L. W. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences, 22*, 577-660.

Barsalou, L. W. (2005). Abstraction as dynamic interpretation in perceptual symbol systems. In L. G.-S. Rakison, *Building object categories*. Mahwah: Erlbaum.

Barsalou, L. W., Santos, A., Simmons, W. K., & Wilson, C. D. (2008). Language and simulation in conceptual processing. In M. D. Vega, A. M. Glenberg, & A. C. Graesser, *Symbols, embodiment, and meaning*. Oxford: Oxford University Press.

Bestgen, Y., & Vincze, N. (2012). Checking and bootstrapping lexical norms by means of word similarity indexes. *Behavior Research Methods, 44*, 998-1006.

Binder, J. R., Desai, R., Graves, W., & Conant, L. L. (2009). Where is the semantic system? A critical review and meta-analysis of 120 functional neuroimaging studies. *Cerebral Cortex*, doi: 10.1093/cercor/bhp055.

Bradley, M. M., & Lang, P. J. (1999). *Affective norms for English words (ANEW): Stimuli, instruction manual and affective ratings (Tech. Rep. No. C-1)*. Gainesville: Center for Research in Psychophysiology, University of Florida.

Bradley, M. M., & Lang, P. J. (2010). *Affective Norms for English Words (ANEW): Affective ratings of words and*

instruction manual (Technical Report C-2). Gainesville: University of Florida.

Foltz, P. W., Kintsch, W., & Landauer, T. K. (1998). The measurement of textual Coherence with Latent Semantic Analysis. *Discourse Processes, 25*, 285-307.

Gallup. (1999, December 31). Mother Teresa voted by American people as most admired person of the century. Retrieved from Gallup News Service: http://www.gallup.com/poll/3367/mother-teresa-voted-american-people-most-admired-person-century.aspx.

Glenberg, A. M. (1997). What memory is for. *Behavorial and Brain Sciences, 20*, 1-55.

Hutchinson, S., & Louwerse, M. M. (in press). Statistical linguistic context and embodiment predict metaphor processing but participant gender determines how much. *Cognitive Linguistics.*

Jones, M. N., Gruenenfelder, T. M., & Recchia, G. (2011). In defense of spatial models of lexical semantics. In L. Carlson, C. Hölscher, & T. Shipley (Eds.), *Proceedings of the 33rd Annual Conference of the Cognitive Science Society* (pp. 3444-3449). Austin, TX: Cognitive Science Society.

Klauer, K. C., & Musch, J. (2001). Does sunshine prime loyal? Affective priming in the naming task. *The Quarterly Journal of Experimental Psychology, Section A, 54*(3), 727-751.

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 Review, 104*(2), 211-240.

Louwerse, M. M. (2008). Embodied relations are encoded in language. *Psychonomic Bulletin & Review, 15*, 838-844.

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

Louwerse, M. M., & Benesh, N. (2012). Representing spatial structure through maps and language: *Lord of the Rings* encodes the spatial structure of Middle Earth. *Cognitive Science, 36*(8), 1556-1569.

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.

Louwerse, M. M., & Jeuniaux, P. (2008). Language comprehension is both embodied and symbolic. In M. d. Vega, A. Glenberg, & A. C. Graesser, *Embodiment and meaning: A debate*. Oxford: Oxford University Press.

Louwerse, M. M., & Jeuniaux, P. (2010). The linguistic and embodied nature of conceptual processing. *Cognition, 114*, 96-104.

Meier, B. P., Hauser, D. J., Robinson, M. D., Friesen, C. K., Schjeldahl, K. (2007). What's "up" with god? Vertical space as a representation of the divine. *Journal of Personality and Social Psychology, 93*, 699-710.

Meier, B. P., & Robinson, M. D. (2004). Why the sunny side is up: Associations between affect and vertical position. *Psychological Science, 15*, 243-247.

Paivio, A. (1978). Mental comparisons involving abstract attributes. *Memory & Cognition, 6*, 199-208.

Paivio, A. (2010). Dual coding theory and the mental lexicon. *The Mental Lexicon, 5*(2), 205-230.

Paivio, A. (2013). Dual coding theory, word abstractness, and emotion: A critical review of Kousta et al. (2011). *Journal of Experimental Psychology: General, 142*(1), 282-287.

Pecher, D., van Dantzig, S., Boot, I., Zanolie, K., Huber, D. E. (2010). Congruency between word position and meaning is caused by task-induced spatial attention. *Frontiers in Psychology, 3*, 1-8.

Pecher, D. & Zwaan, R. A. (Eds.) (2005). *Grounding cognition: The role of perception and action in memory, language, and thinking*. Cambridge: Cambridge University Press.

Pulvermüller, F. (2005). Brain mechanisms linking language and action. *Nature Reviews Neuroscience, 6*(7), 576-582.

Quesada, J. (2006). Creating your own LSA spaces. In T. Landauer, D. McNamara, S. Dennis, & W. Kintsch, *Latent Semantic Analysis: A Road to Meaning*. Mahwah: Erlbaum.

Ranker. (2013, January 9). *The All-Time Worst People in History*. Retrieved from Ranker.com: http://www.ranker.com/crowdranked-list/the-all-time-worst-people-in-history

Semin, G. R., & Smith, E. R. (2002). Interfaces of social psychology with situated and embodied cognition. *Cognitive Systems Research, 3*, 385-396.

Šetić, M., & Domijan, D. (2007). The influence of vertical spatial orientation on property verification. *Language and Cognitive Processes, 22*, 297-312.

Steyvers, M., Shiffrin, R. M., & Nelson, D. L. (2004). Word association spaces for predicting semantic similarity effects in episodic memory. *Experimental Cognitive Psychology and its Applications: Festschrift in Honor of Lyle Bourne, Walter Kintsch, and Thomas Landauer* (pp. 237-249). Boulder, CO: University of Colorado.

Tillman, R., Hutchinson, S., Jordan, S., & Louwerse, M. M. (2013). Verifying properties from different emotions produces switching costs: Evidence for coarse-grained language statistics and fine-grained perceptual simulation. *Proceedings of the 35th Annual Conference of the Cognitive Science Society* (pp. 3551-3556). Austin, TX: Cognitive Science Society.

van Dantzig, S., Zeelenberg, R., & Pecher, D. (2009). Unconstraining theories of embodied cognition. *Journal of Experimental Social Psychology, 45*, 345–351.

Warriner, A. B., Kuperman, V., & Brysbaert, M. (2013). Norms of valence, arousal, and dominance for 13,915 English lemmas. *Behavior Research Methods* , 1-17.

Wikipedia. (2013, January 23). *Wikipedia:Neutral point of view*. Retrieved from Wikipedia: http://en.wikipedia.org/wiki/Wikipedia:Neutral_point_of_view