

Computing Semantics of Preference with a Semantic Cognitive Map of Natural Language: Application to Mood Sensing from Text

Alexei V. Samsonovich and Giorgio A. Ascoli

Krasnow Institute for Advanced Study, George Mason University
4400 University Drive MS 2A1, Fairfax, VA 22030-4444, USA
asamsono@gmu.edu, ascoli@gmu.edu

Abstract

A semantic cognitive map of natural language is constructed using dictionaries of synonyms and antonyms with a simple algorithm described by the authors previously that differs from prior work on embedding of words in metric spaces: Isomap, latent semantic analysis, multidimensional scaling. Previous and new results exhibit semantic invariance across languages and extend to psychometric data. Semantics of the two most significant dimensions can be approximately characterized as “good vs. bad” and “calming vs. exciting”. Applications of this technique in the present work include (i) quantitative definitions of universal semantic dimensions and their experimental validation, and (ii) computation of semantic biases and preferences expressed in an arbitrary given text segment: “mood sensing from text”. The latter, practically important capability can be used in search engines, in human interfaces of intelligent agents, etc. and is illustrated here using a diverse variety of text samples.

Introduction

The two key questions underlying a computational approach to preference making in real life are:

- What metric system or model should be used to represent semantics of preferences (values)?
- How to compute semantics of specific concepts, phrases, etc. within the selected framework?

Both questions are very difficult (e.g., the entire field of economics provides only a partial solution to the problem). Many concepts and their relations are ambiguous and context-dependent; sometimes opposite semantics are intimately connected and switch roles in their evolution. For example, the term “Baroque” was initially used with a derogatory meaning (the word “baroque” in French means “odd”), but later became associated with the respectable cultural epoch. While modern science offers a precise system of measures for physical quantities (SI, the International System of Units), a metric system for human semantic values is missing. Moreover, there is no general consensus on whether it is possible to define a rigorous quantitative model describing the system of human values, and if yes, then what the structure of this model should be.

Copyright © 2008, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Among the most parsimonious possible answers to this question is the notion of a *semantic cognitive map* (SCM: Figure 1) understood as an abstract vector space with symbolic representations allocated as vectors in it, such that synonyms correspond to vectors that are nearly parallel to each other, while antonyms should be nearly anti-parallel (Samsonovich and Ascoli 2007). In this model there are no definite map coordinates a priori assigned to concepts (all coordinates are found by multidimensional optimization of the global energy function: see below). There is no constraint on relative positioning of unrelated concepts. There is no minimal distance among synonyms. All these details clearly distinguish the method of SCM from well-known methods and techniques, including Isomap, multidimensional scaling and latent semantic analysis (see Discussion). Applied to this concept of SCM, the above two questions can be reformulated as follows:

- What are the semantics of the principal dimensions of SCM?
- How to allocate representations of concepts in SCM?

Given answers to these questions, one can think further:

- How to use SCM practically?

These are the questions addressed in the present work.

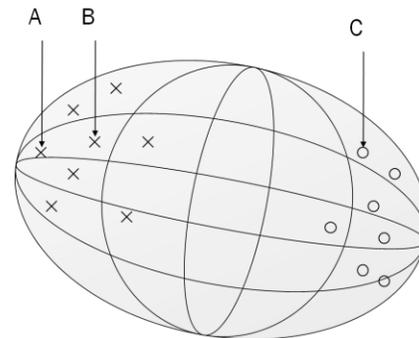


Figure 1. The general concept of a semantic cognitive map. Symbolic representations of concepts (A, B, C) are allocated in an abstract vector space that captures their semantics. A and B are synonyms; A and C are antonyms.

As a specific test material, we take the Microsoft Office 2003-2007 US English dictionary of synonyms and antonyms that includes words and short phrases. At first, the idea seems puzzling: how would one measure the meaning of words without any context? The very idea that

commonsense semantics of words in natural language can be expressed in precise numbers sounds problematic, because of the subjective nature of the attribution of meaning to words. The task of numerical characterization of word meaning has no well-acceptable solution in the literature. E.g., the state of the art in experimental psychology did not change over decades and is based on introspective word ranking by human subjects, e.g., on a scale from 1 to 10 along several fixed semantic dimensions (Osgood et al. 1957, Rubin 1980). Each of these dimensions is selected arbitrarily by the experimenter and is typically given to subjects by a single word characterizing it. Other methodologies involve psychophysical measurements taken during perception of words by human subjects. While results of these experimental studies can be considered as phenomenology, the challenge is to create a theoretical framework that would allow us in principle to express word meaning in numbers precisely. The approach presented here addresses this goal. It should be mentioned that the idea of a possible solution was previously rejected (Tversky and Gati 1982) and then re-vitalized (e.g., Gärdenfors 2004). Results presented here also suggest that a solution is possible.

Semantic Cognitive Map of English

Methods

The term “cognitive map” had been used in cognitive sciences for several decades with various meanings associated with it (Tolman 1948, Downs and Stea 1973, O’Keefe and Nadel 1978). The modern notion introduced by O’Keefe and Nadel (1978) was initially limited to spatial mapping and was subsequently extended to include non-spatial features of contexts. In the present work, a *semantic cognitive map* is understood as a mapping from a set of cognitive representations (e.g. words) to an abstract continuous metric space, such that semantic relations among representations are reflected in geometric relations. Specifically, if words are represented by D -dimensional vectors, then it is expected that vectors of two synonyms will have a positive dot product, while vectors of two antonyms will have a negative dot product. Based on this idea, the energy function H (*) of the system of vectors can be defined as follows (Samsonovich and Ascoli 2007). Given a dictionary of words and phrases and an incidence matrix of their synonym-antonym relations, for each pair of antonyms, the dot product of the corresponding vectors is added to H , and for each pair of synonyms, the dot product of the corresponding vectors is subtracted from H . The last term in H is the sum of fourth powers of all vector length:

$$H = \sum_A xy - \sum_S xy + \frac{1}{4} \sum_x x^4 \quad (*)$$

Here x and y are D -vectors representing words, A is the set of antonym pairs, and S is the set of synonym pairs. The last term reflects the idea that the map should be compact.

Given this definition, a cognitive map is constructed by numerical minimization of H with respect to all vector coordinates, starting from some random set of coordinates. When the minimization is completed, the choice of map coordinates is done based on the *principal components* (PCs) of the final distribution. This method was used in the previous work (Samsonovich and Ascoli 2007) and in the present work, with $D = 26$.

Materials

The material used here is the English Dictionary of synonyms and antonyms available as a part of the thesaurus in Microsoft Word 2003 and 2007 (further referred to as MS Word). The maximal connected component was selected from the graph of synonyms and antonyms restricted to words that have at least two connections, at least one of which is an antonym link. This selection resulted in a set of approximately $N = 16,000$ English words and phrases.

Optimization Results: New vs. Previous Findings

Characteristics of the optimized distribution of words obtained in this study (Figures 2, 3) are consistent with those reported previously (Samsonovich and Ascoli 2007). In particular, the distribution is low-dimensional: only the first four principal components have significant variances.

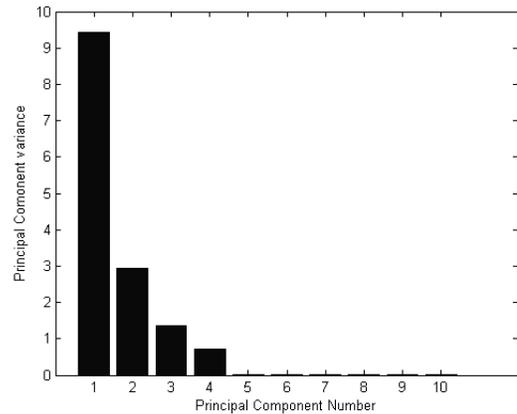


Figure 2. Principal components of the optimized distribution of words. Only the first four principal components out of 26 (only 10 shown) have significant variances. They approximately correspond to the notions of goodness, arousal, freedom and richness.

The main previous finding (Samsonovich and Ascoli 2007) is that map coordinates have definite semantics. E.g., the first principal component corresponds to the notions of “good” and “bad”, and the second principal component corresponds to the notion of “calming” vs. “exciting”. This property was previously validated (Figure 4) by quantitative comparison with the database of Affective Norms for English Words (ANEW) developed by the Center for the Study of Emotion and Attention (CSEA) at the University of Florida. This database contains 1,034 affective English words and was created using the Self-Assessment Manikin to acquire ratings of *pleasure*,

arousal, and *dominance* (Bradley and Lang 1999). In particular, it was found that the first ANEW coordinate, *pleasure*, is the strongest correlate of the first principal component, and the second ANEW coordinate, *arousal*, is the strongest correlate of the second principal component.

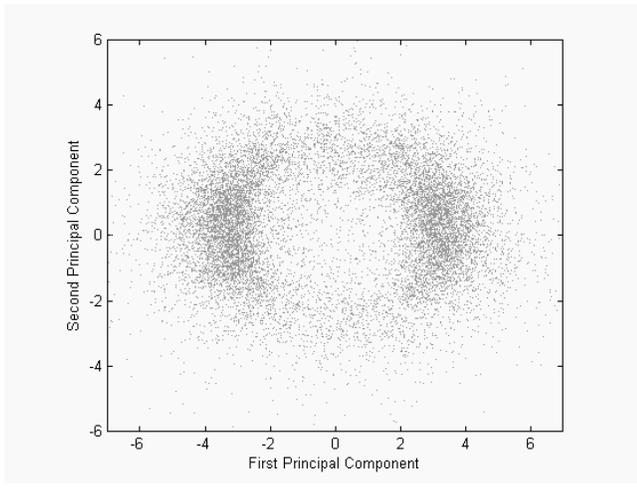


Figure 3. The optimized distribution of 16,000 words in its maximum-spread projection. The shape is determined by the energy function H^* .

Table 1 (reproduced from Samsonovich and Ascoli 2007). Top portions of sorted lists for each PC for two languages. Words that are common across cells within each row are typed in boldface.

PC #	MS Word English	MS Word French: translated automatically
1	increase , well, rise, support, accept , clear, improve, right, continue, direct, good, make, respect, honor, happy , secure, order, understanding, fix , power, bright, present, definite...	happy , agreement, stable, joined together, delighted, approve, net, some, honest, rich, added, increased , pleasant, sincere, union, frank, fix , favor, praise, optimist, accept , abundance, help...
2	calm , easy, soft , gentle, relaxed, light, ease, simple , quiet, soothe, smooth, empty, mild, weak, gently, peaceful, compliant, lenient, pale...	calm , modest, discrete, simple , subjected, thin, alleviated, softened, flexible, sober, moderate, soft , immobility, measured, silence, humble, reserved, simplicity, obeying
3	start, open , fresh, begin, release , original, new, reveal, speed up, free ...	release , deliver, freedom, yield, open , leave, free , disencumbered, discovered, dispersion, broad...

Other previous findings include consistency of results across different languages, as was previously illustrated (Samsonovich and Ascoli 2007) by lists of words sorted along the principal components (Table 1). Moreover, similar semantics (as those observed in the PC#1-PC#2 plane: Figure 3, Table 1) had been used for decades in empirically grounded models of emotions (the *circumplex model*: Russel 1980, Horowitz 2004, Posner et al. 2005). Experimentally identified semantic dimensions of these circumplex models are similar to PC#1 and PC#2 and also exhibit cross-cultural invariance.

The Map Semantics

Given the consistency of findings, the question is: How to identify the observed semantics? The answer is given by coordinates of all words on the map, and there are several possibilities to express it concisely in terms of familiar notions. One is to look at the top or the bottom portions of sorted lists of words (Table 1). Another possibility is to design and then validate a logical expression that captures semantics of principal components, treating words as predicates. For example, the following is a candidate:

$PC\#1 = success \mid positive \mid clear \mid makes-good-sense$

$PC\#2 = exciting \mid sophisticated \mid does-not-go-easy$

An alternative approach is to treat the constructed map and its principal components as definitions of new semantic concepts, called “First Subjective Dimension”, “Second Subjective Dimension”, etc., without attempting to reduce these concepts to a few familiar words or concise definitions based on words. The idea is that these concepts may not be exactly reducible and therefore should not be reduced and replaced with something familiar. These new semantic concepts can be learned by exploration of the map on a case-by-case basis, in analogy with learning natural language. Similarly, the abstract linear unit of the map can be called *linear semantic unit* (LSU: Figure 5).

In order to validate this approach, the following method was used in a pilot study. Random 200 dictionary words were presented to a human participant (one of the authors), and, based on the previous experience obtained during exploration of the map, the participant made a forced decision (without looking at the actual word coordinates) what should be the coordinates of presented words along the First Semantic Dimension. The computed correlation of the subjectively decided coordinate with the actual map coordinate is significantly high (Pearson’s correlation coefficient $r = 0.7$), indicating that the notion of the First Semantic Dimension is unambiguous, can be learned and subsequently used in analysis and processing of linguistic data. Confirmation of these preliminary results with multiple participants would indicate that this notion of First Semantic Dimension can be reliably used in computational measurement of semantics, for communications, etc. Another validation of SCM constructed from the MS Word English dictionary was done previously using the ANEW database (Figure 4: see Samsonovich and Ascoli 2007).

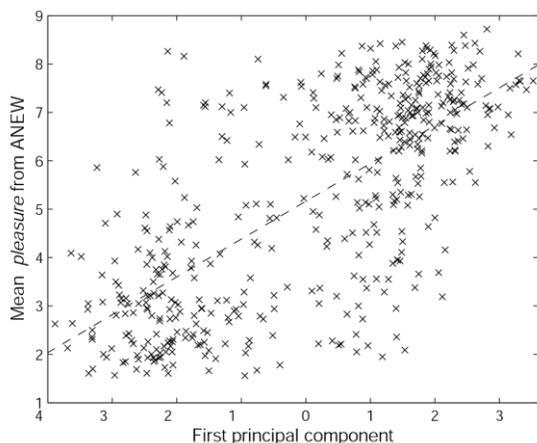


Figure 4 (reproduced from Samsonovich and Ascoli 2007). Scatter plot demonstrating strong correlation of PC #1 with the first dimension of ANEW: *pleasure*. Each cross represents a word. The dashed line is a linear fit. The two clusters formed by “positive” and “negative” words are clearly separated in each map (along each plot dimension).

Application of SCM: Mood Sensing from Text

The semantic cognitive map constructed from MS Word English dictionary can be used to compute semantic preferences expressed in a given text, or to make decision about the mood of a given document. In order to illustrate this possibility, here SCM was used to classify a set of abstracts retrieved from Medline by a search for “neur*” (the search result was truncated at 500). Words from each abstract were located on SCM, and the average vector (the “center of mass”) of all indexed words was computed for each abstract. This is how each abstract was assigned map coordinates. Finally, the set of abstracts was sorted along the first two principal components. Table 2 gives the first two sentences of each of the four “extreme” abstracts.

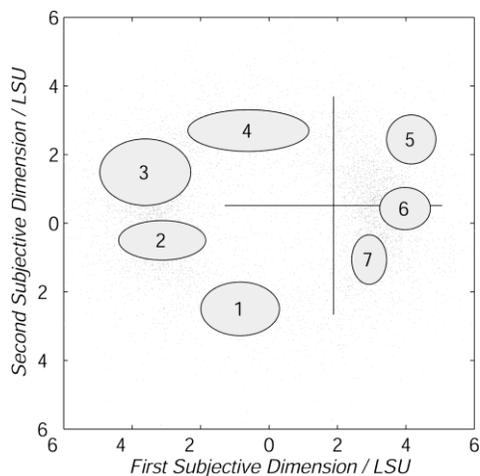


Figure 5. Cognitive map of English with seven utterances automatically allocated on it. The center of each oval is the location of the utterance with the corresponding number, dimensions of the ovals represent standard error, gray dots represent words. LSU: linear semantic unit.

Table 2. The four “extreme abstracts” selected by cognitive mapping out of 500 abstracts retrieved from Medline by a truncated search for “neur*” (only the first two sentences are given for each abstract).

+PC#1: “the most positive”	“The development of noninvasive methods capable of affording ever higher resolution images is an ongoing major objective of scientific investigation for the benefit of both clinical medicine and neuroscience. Since its development in the early 1970s, magnetic resonance imaging (MRI) has remained a technology of choice for medical imaging development because of the wide range of potential clinical applications.”
-PC#1: “the most negative”	“7-Nor-20-oxopregn-5-en-3 beta-yl acetate was converted into (20R)-5 beta,6 beta-epoxy-7-nor-5 beta-pregnane-3 beta,20-diyl diacetate in three steps. Stereospecific migration of the 6 alpha-hydrate ion led to a 6-oxo derivative with a 5 alpha-configuration.”
+PC#2: “the most exciting”	“Drosophila courtship is a complex behavior. A new study shows that glia modulate neurotransmission to influence male preference, but the authors should have resisted the temptation to describe their results in tabloid language.”
-PC#2: “the most calming”	“Tacrolimus is used widely for immunosuppression following transplantation. It has rarely been linked to the development of peripheral neuropathy.”

In order to further validate this method, a set of arbitrarily generated utterances was used. The utterances were intentionally generated with the purpose to cover the entire spectrum of feelings represented by the first two map coordinates, based on a prior experience with exploration of the map and intuitive expectations of the results. These utterances are:

1. *Please, chill out and be quiet. I am bored and want you to relax. Sit back and listen to me.*
2. *Excuse me, sorry, but I cannot follow you and am falling asleep. Can we pause? I've got tired and need a break.*
3. *I hate you, stupid idiot! You irritate me! Get disappeared, or I will hit you!*
4. *What you are telling me is terrible. I am very upset and curious: what's next?*
5. *Wow, this is really exciting! You are very smart and brilliant, aren't you?*
6. *I like very much every word that you say. Please, please, continue. I feel like I am falling in love with you.*
7. *We have finally found the solution. It looks easy after we found it. I feel completely satisfied and free to go home.*

The results of automated allocation of these utterances on SCM are represented in Figure 5. The plane of the figure corresponds to the first two semantic coordinates, labeled here “First Subjective Dimension” and “Second Subjective Dimension”, respectively. The fuzzy annulus of grey dots in the background are the dictionary words allocated on the map (same as in Figure 3). The center of the cross in the middle, shifted to the right, represents the average of all word vectors weighted by the general frequency of usage of words. The shaded ovals represent estimates of feelings associated with the given sample: the center of an oval is the estimate, and the dimensions of each oval represent the standard error. Ovals are numbered according to the numbers of utterances in the list above. The significant dimensionality of the sample is equal 2 (calculated by MANOVA, $P < 0.003$).

A similar mapping study was performed with other samples of text: (i) a set of 500 Medline abstracts, as described above (Table 2), (ii) the set of 152 Shakespeare sonnets, (iii) a set of 10 abstracts of accepted grants, (iv) a set of 45 anecdotes about famous personalities. All results show qualitative consistency with each other.

Discussion: Related Works and Techniques

The method of semantic cognitive mapping described here and in our previous work (Samsonovich and Ascoli 2007) does not stand aside from the mainstream research, and needs to be placed in the relevant context. As mentioned above, methods of geometry and topology are popular in data analysis concerned with semantic dissimilarity. Here the two main building blocks are (a) geometric representation of semantics of individual data items themselves and (b) geometric representation of semantic relations. Accordingly, the two general problems are (1) getting from (a) to (b) and (2) getting from (b) to (a). Popular approaches to (1) include latent semantic analysis and cluster analysis. As to the second problem, one, a rather nonstandard, approach to it is based on SCM (Samsonovich and Ascoli 2007).

The general problem (2) can be stated as follows: given a matrix W of mutual distances or dissimilarities of n data items, compute their best embedding as points in an abstract space, given by the coordinates of points Y , that minimizes a certain criterion $H(Y, W)$: the energy/cost function, which in the literature is also called the loss function. Two best-known techniques used in connection with (2) are multidimensional scaling (MDS: Seber 1984, Borg and Groenen 2005) and Isomap, that relies on MDS (Tenenbaum 2000). In order to explain how these methods differ from our approach and why they should produce different results from ours, we need to briefly outline them.

Classical metric multidimensional scaling procedure takes a full matrix W of distances among n points and computes an n -by- n matrix Y of n points embedded in an n -dimensional Euclidean space \mathbf{R}^n . Simultaneously, the eigenvalues e of Y^*Y' are computed. If the distances W are Euclidean, and, e.g., the first D elements of e are much

larger than the rest $n-D$ elements, then the result Y is truncated to the first D columns (dimension reduction). In other words, n points are embedded in a smallest space \mathbf{R}^D , such that their mutual distances W are (approximately) preserved. When W is a more general dissimilarity matrix that does not satisfy the metrics axioms (Tversky and Gati 1982), then some elements of e may be negative, and then D is determined as the number of positive eigenvalues, providing an embedding that gives a reasonable approximation to W , if the negative elements of e are small in magnitude (Seber 1984).

While classical multidimensional scaling requires specification of distances (or dissimilarities) for each pair of points, non-classical metric multidimensional scaling procedure allows for missing distance data, and offers to limit the embedding space dimension by an ad hoc number that is lower than D (in this case, an additional criterion for selecting the best embedding needs to be specified).

Nonmetric multidimensional scaling procedure creates a configuration of points whose mutual distances approximate the given dissimilarities W ; however, instead of trying to interpret given dissimilarities themselves as distances among points, non-metric multidimensional scaling finds a configuration in which distances among points are equal to some nonlinear monotonic function of W . The unspecified nonlinearity of this function implies that this procedure only attempts to preserve the ordering of dissimilarities, essentially ignoring their values.

Isomap (Tenenbaum 2000) is similar to nonclassical multidimensional scaling in that it allows for missing, in fact, very sparse dissimilarity data. The idea of handling the problem (2) that underlies Isomap is to compute each missing matrix element as the length of the shortest path in the graph representing dissimilarities, in which lengths of links are proportional to the matrix elements of W . Finally, a classical multidimensional scaling procedure is applied.

The outlined above and similar techniques are suitable for handling one semantic relation that can be reduced to a semantic distance. In contrast, the approach presented in this work requires two relations: synonymy and antonymy that are not reducible to each other (antonyms on average are more closely related to each other than unrelated words). This detail is significant: if one of the two relations is disregarded, then SCM produces dramatically different results (Samsonovich and Sherrill 2007). Plus, the idea of SCM is to optimize the sum of dot products (*) rather than to match given distances. The criterion (*) does not penalize unrelated concepts for being close to each other.

All these details differentiate our SCM from the well-known methods outlined above. Therefore, the results of optimization are different in both cases: e.g., the dimension of the resultant embedding should be lower in SCM than it can be expected when Isomap or classical MDS is used.

The presented approach is also very different from modern trends in NLP research aimed at automated mood classification using a combination of syntactic and semantic parsing (e.g., Jung et al. 2006, Wiebe et al. 2005). Generally, the problem is considered extremely difficult by

expert computational linguists (Dennis Perzsanowski, personal communication). The strength of the SCM-based approach is in the availability of quantitative semantic information virtually for all dictionary words rather than for a small subset of words (e.g., as in the study of Jung et al. 2006 that used ANEW). The weakness of the method presented here is that all words are treated equally regardless of available syntactic information. Therefore, combining the parsing approaches with SCM can be advantageous to both and may enable new capabilities.

Concluding Remarks

The main conclusion for this study is that the cognitive mapping method presented here introduces new, universal computational semantic measures that apply to natural language. These measures in principle can be validated and used for many purposes, e.g., for automated detection of mood and semantic preferences expressed in an arbitrary coherent text segment, thus opening new perspectives for NLP related to human-computer interface and data mining.

In a broader artificial intelligence context, SCM similar to the one constructed here can be used to give human-like “feelings” and “affects” to artifacts. For example, suppose that a dynamical working memory state of an agent is characterized by a set of active symbolic representations, each of which is associated with a word or a phrase from natural language. Then, the method described here can be used to compute the current mood of the agent based on the set of words associated with its current state.

In general, in order to achieve a human level of emotional intelligence in an artifact, it will be probably necessary to use an SCM derived from the entire human culture. Constructing this SCM is a big and significant challenge, comparable to the human genome challenge. The presented study does not solve the problem, but it indicates that the problem is solvable. Specifically, we demonstrated a possibility to define universal, computable semantic characteristics that can provide a basis for the autonomous development of a system of values in the future artificial cognitive systems.

Acknowledgments

The authors are grateful to Dr. Kenneth A. De Jong for fruitful discussions of preliminary results.

References

Borg, I. and Groenen, P., 2005. *Modern Multidimensional Scaling: Theory and Applications, Second Edition*. New York: Springer-Verlag.

Bradley, M. M., and Lang, P. J. 1999. Affective norms for English words (ANEW): Stimuli, instruction manual and affective ratings. *Technical report C-1*. Gainesville, FL: University of Florida.

Downs, R. M., and Stea, D. 1973. Cognitive maps and spatial behavior: Process and products. In Downs, R. M., and Stea, D. (Eds.). *Image and Environments*, pp. 8-26. Chicago: Aldine.

Gärdenfors, P. 2004. *Conceptual Spaces: The Geometry of Thought*. Cambridge, MA: MIT Press.

Horowitz, L. M. 2004. *Interpersonal foundations of psychopathology*. Washington, DC: American Psychological Association.

Jung, Y., Park, H., and Myaeng, S. H. 2006. A hybrid mood classification approach for blog text. *Lecture Notes in Computer Science* 4099: 1099-1103.

O’Keefe, J., and Nadel, L. 1978. *The Hippocampus as a Cognitive Map*. New York, NY: Clarendon.

Osgood, C.E., Suci, G., and Tannenbaum, P. 1957. *The measurement of meaning*. Urbana, IL: University of Illinois Press.

Posner, J., Russell, J. A., and Peterson, B. S. 2005. The circumplex model of affect: An integrative approach to affective neuroscience, cognitive development, and psychopathology. *Development and Psychopathology* 17: 715-734.

Rubin, D. C. 1980. 51 properties of 125 words: a unit analysis of verbal behavior. *Journal of Verbal Learning and Verbal Behavior* 19: 736-755.

Russell, J. A. 1980. A circumplex model of affect. *Journal of Personality and Social Psychology* 39 (6): 1161-1178.

Samsonovich, A. V. and Ascoli, G. A. 2007. Cognitive map dimensions of the human value system extracted from natural language. In Goertzel, B. and Wang, P. (Eds.). *Advances in Artificial General Intelligence: Concepts, Architectures and Algorithms. Proceedings of the AGI Workshop 2006. Frontiers in Artificial Intelligence and Applications*, pp. 111-124. IOS Press: Amsterdam, The Netherlands.

Samsonovich, A. V. and Sherrill, C. P. 2007. Comparative study of selforganizing semantic cognitive maps derived from natural language. In D. S. McNamara & J. G. Trafton (Eds.), *Proceedings of the 29th Annual Cognitive Science Society*, p. 1848. Austin, TX: Cognitive Science Society.

Seber, G. A. F., 1984. *Multivariate Observations*. Wiley.

Tenenbaum, J. B. 2000. A global geometric framework for nonlinear dimensionality reduction. *Science* 290: 2319.

Tolman, E. C. 1948. Cognitive maps in rats and man. *Psychological Review* 55 (4): 189-208.

Tversky, A. and I. Gati 1982. Similarity, separability, and the triangle inequality. *Psychological Review* 89(2): 123-154.

Wiebe, J., Wilson, T., and Cardie, C. 2005. Annotating expressions of opinions and emotions in language. *Language Resources and Evaluation*, 39 (2-3): 165-210.