Concept similarity: An abstract relevance classes approach

Luigi Lombardi (luigi.lombardi@unitn.it)
Department of Cognitive Sciences and Education, University of Trento
I-36068 Rovereto (TN), Italy

Giuseppe Sartori (giuseppe.sartori@unipd.it)
Department of General Psychology, University of Padua
I-35100 Padova, Italy

Abstract

In this paper we present a new model, called ARCLASS, for computing similarity between concepts. ARCLASS generalizes the Tversky's Ratio model of similarity (Tversky, 1977) and uses a non-negative matrix factorization procedure to estimate abstract classes on which similarity between concepts may be computed. We applied the model to a high-dimensional database of semantic features used to describe Living and Non-living concepts. Our results suggest that similarity between concepts may be at the origin of many disorders of conceptual knowledge and, therefore, ARCLASS may be considered as a valuable measurement tool to test hypotheses about category-specific disorders.

Introduction

Modeling the capacity to "judge" one stimulus or object as similar to another is necessary for understanding many cognitive processes, including perception, categorization, and action. However, similarity is still a matter of debate and a number of empirical and theoretical arguments have undermined its role and validity in cognitive processing (Medin, Goldstone, & Gentner, 1993; Goldstone & Medin, 1994). Formal models of similarity can be distinguished into two broad classes: geometrical models (e.g., Nosofsky, 1986; Shepard, 1987; Kruschke, 1992) and featural models, also called set-theoretical models (e.g., Tversky, 1977). The geometrical models represent stimuli as points in a multidimensional metric (typically Euclidean) space and similarity is treated as a decreasing function of distance. In contrast, featural models represent stimuli as salient subsets of features with similarity defined as a function of the subsets of common and distinctive features. More specifically, similarity increases as the number of common features between objects increases, and as the number of features possessed by only one object (distinctive features) decreases.

Some authors (e.g., Tversky, 1977) observed that geometrical models are better suited to domains where stimuli vary continuously along a relatively small number of dimensions, whereas the discrete nature of the featural models makes them more appropriate for modeling domains where stimuli are defined in terms of a large number of properties or features. Both approaches provided a solid foundation for the similarity assessment procedures used in explaining data from psychological experiments of human similarity judgement. Moreover, recently some new interesting model extensions have been proposed within the two approaches (e.g., Tenenbaum & Griffiths, 2001; Love, Medin, & Gureckis, 2004; Navarro & Lee, 2004).

In this paper we will present a new similarity model, called ARCLASS (Abstract Relevance Classes Modeling), for the computation of concept similarity in human concept representations. ARCLASS immediately links up with previous models developed within the featural approach and is based on the notion of *semantic relevance* (Sartori & Lombardi, 2004). Semantic relevance (SR) is a continuous parameter which represents the contribution of a feature to the cognitive representation "core" of a concept. More precisely, in a cognitive system, the notion of relevance of features is intended to capture the property of a given feature in distinguishing one concept from other similar ones. SR has been shown to play a significant role in the modeling of semantic memory of normal or impaired cognitive systems (Sartori & Lombardi, 2004; Mechelli, Sartori, Orlandi, & Price, 2005; Sartori, Lombardi, & Mattiuzzi, 2005; Sartori, Polezzi, Mameli, & Lombardi, 2005).

Technically, ARCLASS represents an extension of the well known Tversky's (1977) Ratio Model of similarity and uses a non-negative matrix factorization (NMF) procedure (Lee & Seung, 1999) to estimate few basic abstract relevance classes from a high-dimensional featural representation. The latter procedure yields a low-dimensional semantic representation from which concept similarities may be computed.

The paper is organized as follows. First, we introduce the main formal aspects of the new model. Next, we proceed with reporting an application of this new model to real high-dimensional data on concept similarity. Finally, our conclusions are given in the last section.

The ARCLASS framework

Semantic relevance

Within a featural representation, concepts are defined by vectors of weights codifying the intensities of features or properties used in describing a concept domain. The concept domain $\mathcal{D}$ is described by a finite set of I different concepts (c_i) and a finite set of J different features (f_j), respectively. It is worthwhile to represent $\mathcal{D}$ as an $I \times J$ intensity matrix $\mathbf{X} = [x_{ij}]$, where $x_{ij} \in \mathbb{R}_+$ denotes the degree of association (*intensity*) between Feature j and Concept i. If $x_{ij} = 0$, we say that c_i and f_j are unrelated in $\mathcal{D}$.

Formally, a relevance process acts by transforming $\mathbf{X}$ into an $I \times J$ relevance model matrix $\mathbf{K} = [k_{ij}]$, which represents the relevance model for $\mathcal{D}$. The fundamental assumption of a relevance model is that $\mathbf{K}$ can be decomposed into an $I \times J$ matrix $\mathbf{L}$ and a $J \times J$ diagonal matrix $\mathbf{G}$, by means of the matrix product:

$$\mathbf{K} = \mathbf{L}\mathbf{G} \tag{1}$$

In the above equation, $\mathbf{L}$ represents an $I \times J$ matrix of weights with entry l_{ij} of $\mathbf{L}$, denoting the local importance of Feature j for Concept i; hence, $\mathbf{L}$ is called the local importance matrix. Main diagonal $diag(\mathbf{G})$ of $\mathbf{G}$ represents a vector of J weights with entry g_j of $diag(\mathbf{G})$, denoting the overall importance of Feature j for all I concepts; hence, $\mathbf{G}$ is called the global importance matrix.

$\mathbf{L}$ and $diag(\mathbf{G})$ may be derived by means of two weighting mappings $\langle \phi, \psi \rangle$:

$$\mathbf{L} = \phi(\mathbf{X}), \qquad diag(\mathbf{G}) = \psi(\mathbf{X}) \tag{2}$$

which act as a linking structure between intensity matrix $\mathbf{X}$ and relevance matrix $\mathbf{K}$.

Several weighting schemes may be derived to model relevance (Sartori & Lombardi, 2004; Sartori et al., 2005). In this paper, we refer to a simple weighting scheme called FI $\times$ ICI (*Feature Intensity* $\times$ *Inverse Concept Intensity*), adapted from Salton's well-known TF $\times$ IDF (*Term Frequency* $\times$ *Inverse Document Frequency*) information retrieval measure (Salton, 1989). Under the FI $\times$ ICI assumption, we set:

$$l_{ij} = \phi(x_{ij}) = x_{ij} \tag{3}$$

$$g_j = \psi(\mathbf{x}_{\cdot j}) = \log_2\left(\frac{I}{I_j}\right) \tag{4}$$

$(\forall i = 1, \ldots, I; \forall j = 1, \ldots, J)$ with $\mathbf{x}_{\cdot j}$ and $0 < I_j \leq I$ respectively denoting the j^{th}–column of $\mathbf{X}$ and the number of concepts in which Feature j loads a positive intensity (that is for which $x_{ij} > 0$). From the above conditions it follows that

$$k_{ij} = x_{ij} \times \log_2\left(\frac{I}{I_j}\right). \tag{5}$$

Therefore, a feature which captures the core of the cognitive representation of a concept will have both high local importance and high global importance.

The mappings ϕ and ψ may be given substantive interpretations. Mapping ϕ is a local weighting function. It is a measure of the strength of a feature in describing a concept; the higher is the intensity x_{ij} of Feature f_j for Concept c_i, the more it is dominant for that concept (Ashcraft, 1978). Note that, under the FI $\times$ ICI assumption, mapping ϕ reduces to the identity function. Mapping ψ is a global weighting function. It is an inverse function of the degree of sharedness of f_j within the concept set. In other words, the more concepts a given feature is connected with, the less distinctive is that feature[1]. The assumption justifying the usage of this mapping is that distinctive features would be overall more important than would shared features, as they are more informative in distinguishing one concept from others (Marques, 2005; Tversky, 1977).

Abstract relevance classes

Given that cognitive operations are limited in their capacity, it is plausible that similarity judgments depend on reducing the complexity of information to a level that does not exceed capacity. Therefore, we assume that concepts can also be described by unobservable features

$$\alpha_1, \ldots, \alpha_M, \quad \text{where} \quad M \ll \min\{I, J\}$$

called *abstract relevance classes*. The reason for assuming the existence of abstract relevance classes is twofold. First, under limited cognitive resources, it does not seem reasonable to assume that a cognitive system uses the whole relevance information to compute concept similarity (Elman, 1993; Halford, Baker, McCredden, & Bain, 2005). Second, abstract relevance classes constitute natural superordinate information that the system can utilize to cluster concepts and, therefore, boost similarity evaluations.

Given the initial relevance model matrix $\mathbf{K}$ (previously derived from $\mathbf{X}$ using the weighting mappings $\langle \phi, \psi \rangle$), the abstract relevance classes problem is to find two new reduced-dimensional matrices $\mathbf{Y}$ $(I \times M)$ and $\mathbf{Z}$ $(J \times M)$, to approximate $\mathbf{K}$ by the product $\mathbf{K}^* = \mathbf{Y}\mathbf{Z}'$ in terms of some loss function. In the approximation of $\mathbf{K}$, M denotes the rank of the model, that is to say, the number of abstract classes adopted in the model. $\mathbf{Y}$ includes M column vectors, called abstract concept relevance bundles, and hence $\mathbf{Y}$ is called the abstract concept matrix. Similarly, $\mathbf{Z}$ includes M column vectors, called abstract feature relevance bundles, and hence is called the abstract feature matrix.

Entry y_{im} of $\mathbf{Y}$ may be interpreted as the abstract relevance of α_m for Concept i, whereas, entry z_{jm} of $\mathbf{Z}$ denotes the activation of Feature j given class α_m. The original relevance value k_{ij} is approximated by the sum:

$$k_{ij} \approx k^*_{ij} = \sum_{m=1}^{M} y_{im} z_{jm}. \tag{6}$$

In other words, the relevance k_{ij} of Feature j for Concept i is approximated by the sum of the M abstract relevances for Concept i weighted by the corresponding activations of Feature j. The approximation of $\mathbf{K}$ is done such that for a fixed rank M the loss function

$$L_2(\mathbf{Y}, \mathbf{Z}) = \left(\sum_{i=1}^{I}\sum_{j=1}^{J}(k_{ij} - k^*_{ij})^2\right)^{\frac{1}{2}} \tag{7}$$

[1]We assume that all J features are singly informative for at least one concept. That is to say, each feature is connected with at least one concept. From the latter, it follows that a feature j shows maximal distinctiveness when connected with one concept only, that is, when $I_j = 1$. Dually, a feature j has minimal distinctiveness when feature j is connected with all the I concepts, that is to say, when $I_j = I$).

is minimized subject to $y_{im} \geq 0$ and $z_{jm} \geq 0$ for each i,j and m. Notice that the inherent nonnegative representation in $\mathbf{K}$ is preserved by the matrix decomposition as the result of the constraints placed on factorization $\mathbf{K}^* = \mathbf{Y}\mathbf{Z}'$ that produce nonnegative lower rank factors that can be interpreted as semantic bundles in the concept domain. The concept vectors in the original relevance matrix $\mathbf{K}$ can be reconstructed by combining these semantic bundles and, therefore, concept similarity can be obtained by comparing the resulting row patterns in the abstract concept matrix $\mathbf{Y}$.

Algorithm

Algorithms designed to approximate $\mathbf{K}$ by solving the constrained minimization problem (Eq. (7)) generally begin from an initial random configuration for $\mathbf{Y}$ and $\mathbf{Z}$. The random assignment is constrained to non-negative values. The routine then follows by alternating iterations to improve the estimates of $\mathbf{Y}$ and $\mathbf{Z}$. The factorization algorithm used in this paper consists in the multiplicative update rules of Lee and Seung (2001). One of the advantage of the multiplicative update rules is that the loss function L_2 is monotonically nonincreasing and becomes constant if and only if $\mathbf{Y}$ and $\mathbf{Z}$ are at a stationary point of L_2 (Lee & Seung, 2001).

Similarity matching

Also concept similarities can be represented in a matrix format $\mathbf{P} = [p_{ih}]$, where $\mathbf{P}$ is a square $(I \times I)$ possibly asymmetric matrix which takes values in $[0, 1]$. Entry p_{ih} of $\mathbf{P}$ denotes the degree of similarity between Concept i and Concept h. If $p_{ih} = 0$, we say that c_i and c_h are fully distinguishable, whereas, if $p_{ih} = 1$, then c_i and c_h are totally undistinguishable.

In ARCLASS the similarity of two concepts p_{ih} is modelled by a generalization of the Tversky's Ratio model (Tversky, 1977). In particular, the similarity between Concept i and Concept h is defined as a weighted additive measure of their common and distinctive abstract relevance classes which takes the following form:

$$p_{ih} = \frac{\sum(\mathbf{y}_i \cap \mathbf{y}_h)}{\sum(\mathbf{y}_i \cap \mathbf{y}_h) + \beta \sum(\mathbf{y}_i - \mathbf{y}_h) + \gamma \sum(\mathbf{y}_h - \mathbf{y}_i)}, \tag{8}$$

where $\mathbf{y}_i$ (resp. $\mathbf{y}_h$) denotes the i^{th}-row (resp. h^{th}-row) of the abstract concept matrix $\mathbf{Y}$, and where vector operations intersection $(\cap)$ and difference $(-)$ are defined as

$$\begin{aligned}
\mathbf{y}_i \cap \mathbf{y}_h &= (\min\{y_{im}, y_{hm}\} : m = 1, \ldots, M), \\
\mathbf{y}_i - \mathbf{y}_h &= (\max\{0, y_{im} - y_{hm}\} : m = 1, \ldots, M), \\
\mathbf{y}_h - \mathbf{y}_i &= (\max\{0, y_{hm} - y_{im}\} : m = 1, \ldots, M).
\end{aligned}$$

Finally, $\sum(\cdot)$ and parameters $\beta, \gamma \geq 0$, respectively, represent the summation function

$$\sum(\mathbf{c}) = \sum_{m=1}^{M} c_m, \quad \mathbf{c} \in \mathbb{R}_+^M,$$

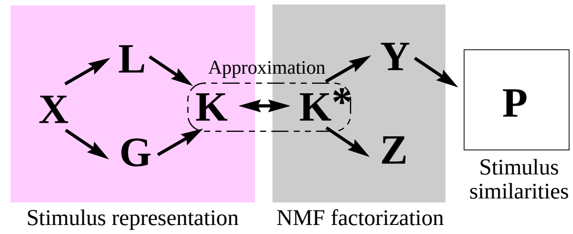

Figure 1: Step by step procedure to compute concept similarities.

and different degrees of importance of distinctive components $\mathbf{y}_i - \mathbf{y}_h$ and $\mathbf{y}_h - \mathbf{y}_i$. Notice that, $p_{ih} = 0$ if and only if $\mathbf{y}_i \cap \mathbf{y}_h = \mathbf{0}$, whereas $p_{ih} = 1$ if and only if $\mathbf{y}_i = \mathbf{y}_h$. An overall picture of the entire computation process is illustrated in Figure 1.

An empirical application

In this section we will illustrate the new approach with an application in the field of semantic memory. In particular, we focus on the evaluation of concept similarity within natural language categories in semantic memory. ARCLASS was applied to a database of cued verbal descriptions of 254 concepts belonging to a corpus collected by Dell'Acqua, Lotto and Job (2000). A total of 2619 features were extracted from normative verbal descriptions (for details see Sartori and Lombardi, 2004) yielding a 254 (concepts) $\times$ 2619 (semantic features) intensity matrix $\mathbf{X}$. Entry x_{ij} of $\mathbf{X}$ was set equal to the number of co-occurrences of Feature j in Concept i over all subjects' descriptions. Finally, relevance weights for semantic features were computed using Eq. (5).

Similarity in low rank reduction.

The aim of this section is to empirically investigate how concept similarity in human subjects is modulated by the number of abstract classes involved in the factorization model. The problem of how to quantify human information processing capacity is considered crucial in cognitive modelling of semantic memory processes. An optimal process depends on reducing the complexity of information that has to be processed so that the amount of cognitive costs does not exceed processing capacity (Elman 1993; Miller 1956). Our main hypothesis is that cognitive systems integrate a limited number of abstract features to compute concept similarity. An evaluation of the low-dimensionality effect on concept similarity can be obtained by comparing the similarities computed on the basis of few abstract features with those processed given the whole set of observable features. Notice that, in terms of knowledge representation, the latter condition reduces to an optimal scenario in which the whole information is taken into account in similarity matching. Therefore, the loss in similarity performance due to dimensionality reduction may provide a direct measure of the suboptimal processing behavior in semantic memory systems.

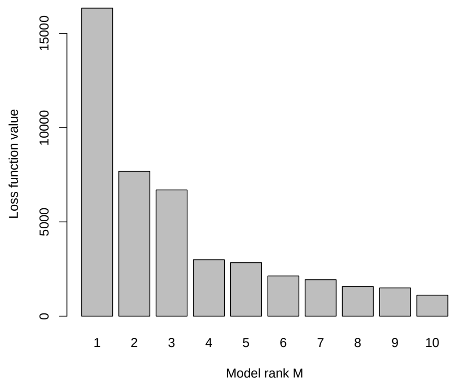

Figure 2: Loss function value $E(\mathbf{P}^\star, \mathbf{P}_M)$ as a function of the model rank $M = 1, 2, \ldots, 10$.

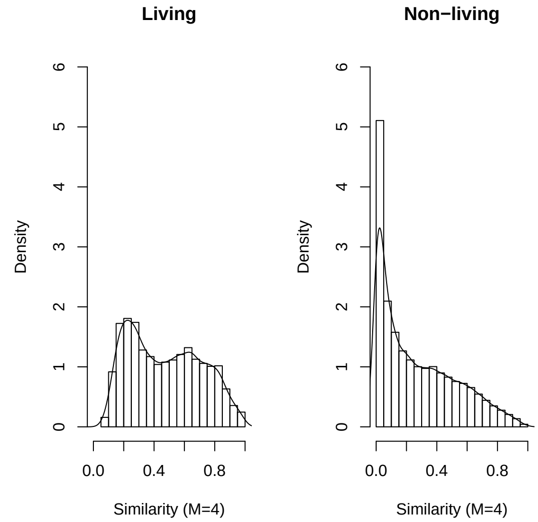

Figure 3: Category histograms based on $\mathbf{P}_4$ reduction. Left side of the picture represents Living concepts. Right side of the picture shows Non-living concepts.

In order to evaluate this aspect, we ran the ARCLASS algorithm in ranks 1-10 using the original relevance model matrix $\mathbf{K}$ as model input[2]. The latter yielded ten approximated relevance matrices

$$\mathbf{K}^*_M = \mathbf{Y}_M \mathbf{Z}'_M,$$

one for each model rank M ($M = 1, \ldots, 10$). Next, the associated similarity matrices $\mathbf{P}_M$ were computed from the abstract concept matrices $\mathbf{Y}_M$ using Eq. (8) with similarity parameters $\beta = \gamma = 1$. Finally, in order to evaluate the optimal concept similarity representation, we also computed the similarity matrix $\mathbf{P}^\star$ containing the set of all pairwise concept similarities calculated on the basis of the total $J = 2619$ observable features. The latter matrix was obtained by replacing $\mathbf{Y}_M$ with $\mathbf{K}$ in Eq. (8). In order to evaluate the suboptimal representation of concept similarity, as a consequence of low rank model reductions, we used loss function (see also Fig. 2)

$$L(\mathbf{P}^\star, \mathbf{P}_M) = \sum_{i=1}^{I} \sum_{h=1}^{I} (p^*_{ih} - p^M_{ih})^2 \qquad (9)$$

Note that in Figure 2 an elbow is present at $M = 4$ which suggests that four abstract classes might be appropriate in reconstructing with sufficient accuracy ($r = 0.68$) the concept similarity representation. On the basis of the above mentioned results, hereafter, we decided to continue the analysis on the four classes reduction $\mathbf{P}_4$.

Similarity distributions of Living and Non-living categories. We analyzed similarity distributions[3] within Living concepts (LV, $N = 96$) and Non-living concepts (NLV, $N = 156$), where N refers to sample size. Furthermore, we also investigated the similarity distributions of four Living subcategories: animals (ANI, $N = 43$), fruits (FRU, $N = 20$), plants (PLA, $N = 12$), vegetables (VEG, $N = 21$), as well as of four Non-living subcategories: clothes (CLO, $N = 19$), kitchen utensils (KU, $N = 19$), transportation vehicles (TV, $N = 23$), musical instruments (MI, $N = 14$). In statistical terms a category $C1$ shows a larger within-similarity than another category $C2$ if and only if the distribution of $C1$ is stochastically larger than that of $C2$. In order to check dominance relations among the concept categories we adopted the Wilcoxon-Mann-Whitney test (see Siegel & Castellan, 1988) which is a non-parametric statistic well suited for highly skewed distributions. The similarity density distributions for the two superordinate categories are shown in Figure 3. Overall LV concepts were more similar to one another than NLV ones ($p < 0.001$).

A Living (resp. Non-living) interval order (Fishburn, 1970) can be obtained by analyzing all pairwise subcategory dominance relations. The LV interval order was as follows:

$$(\text{FRU} \sim (\text{VEG} \succ^* \text{PLA})) \succ^* \text{ANI},$$

where $\sim$ and $\succ^*$ denote stochastic equivalence and stochastic dominance[4], respectively. Similarly, for Non-living subcategories we observed the following interval order

$$(\text{CLO} \sim \text{MI}) \succ^* \text{KU} \succ^* \text{TV}.$$

The distributional within-category similarities, as resulting from the application of ARCLASS, can be used

[2] We remind the reader that in an ARCLASS model the rank M denotes the number of abstract relevance classes represented in the model.

[3] Given $\beta = \gamma = 1$, in ARCLASS a similarity distribution of N different concepts contains all $[(N-1)N]/2$ pairwise similarities.

[4] The superscript * stands for *statistically significant* at $\alpha = 0.001$. We corrected all p-values using the false discovery rate statistic for multiple comparisons (Benjamini & Hochberg 1995).

to highlight important controversial issues about impairments of semantic memory. Cognitive neuroscientists have shown how neurological patients with semantic memory disorders may show selective impairments for certain categories but not for others. In particular, it has been suggested that high similarity between exemplars of the same category may reduce the retrieval accuracy in patients suffering from semantic disorders (Humphreys & Forde, 2001). For example, given that Musical instruments are more confusable than other Non-living concepts this may explain why some patients who are spared on Non-living may be impaired on Musical Instruments (Dixon, Piskopos, & Schweizer, 2001). Several consistent trends in terms of the categories that tend to be impaired/spared together have been reported within the neuropsychological literature (Capitani, Laicona, Mahon, & Caramazza, 2003). The following are established facts in semantically impaired patients: a) Living concepts deficits are much more frequent than Non-living concepts deficits, b) Fruits/vegetables pattern together and can be separately impaired, c) Non-living categories pattern together; these exclude musical instruments that can be impaired along with living concepts. The highlighted within-category similarity clearly parallels the empirical results on semantically impaired patients as reported above.

Interpreting abstract relevance classes and concept clustering. A further support for the plausibility of ARCLASS as a model for semantic similarity may be obtained by analyzing the meaning of the abstract relevance classes in the model decomposition. Meaningful interpretations of the four abstract classes $\alpha_1, \alpha_2, \alpha_3$ and α_4 can be obtained by analyzing the top-most activated concepts in the abstract stimulus matrix $\mathbf{Y}_4$. As an example, Table 1 reports in descending order (for each abstract class), the ten most relevant concepts. In substantive terms the abstract stimulus matrix $\mathbf{Y}_4$ reads as follows: The grouping of concepts connected with α_1 is in line with the representation of Non-living concepts.

In particular, α_1 clusters together clothes (top-most ranking [1-18]) and tools/furnitures (middle ranking [19-40]). Likewise α_2 groups together mostly Non-living concepts (eg. transportation vehicles, buildings, tools), although the ranking does not reflect sharp separable clusters. Class α_3 splits musical instruments (top-most ranking [1-15]) from tools/weapons. Finally, α_4 reflects an hidden dimension codifying Living concepts: fruits (approximately top-most ranking [1-11]) and vegetables/flowers (middle ranking). Interestingly, although fruits and vegetables pattern together they can still be separately impaired as they occupy different ranks in the abstract class (fruits: top-most ranking; vegetables: middle-ranking). Finally, also animals were clustered together (ranks [50-100]). The latter results suggest that similarity between concepts may be at the origin of many disorders of conceptual knowledge and, therefore, ARCLASS may be considered as a valuable measurement tool to test hypotheses about category-specific disorders.

Table 1: Ten best concept examples for each abstract class α_m (ordered by decreasing abstract relevance). Underlined items refer to Living concepts.

Rank	α_1	α_2
1	(9.40) Sweater	(8.56) Locomotive
2	(9.06) Dress	(8.12) Train
3	(8.25) Shirt	(7.57) Ambulance
4	(8.14) Undershirt	(5.46) Trolley
5	(8.13) Pants	(4.96) Car
6	(8.02) Jacket	(4.86) Boat
7	(7.64) Shoes	(4.65) Truck
8	(7.56) Socks	(4.12) Bicycle
9	(6.89) Boots	(3.86) Carriage
10	(6.48) Necktie	(3.22) Motorboat
Rank	**α_3**	**α_4**
1	(8.13) Guitar	(10.04) <u>Lemon</u>
2	(8.09) Piano	(8.86) <u>Apricot</u>
3	(7.42) Accordion	(8.58) <u>Orange</u>
4	(7.28) Bagpipes	(8.35) <u>Apple</u>
5	(6.63) Trumpet	(8.29) <u>Pineapple</u>
6	(6.48) Saxophone	(7.67) <u>Fig</u>
7	(6.08) Flute	(7.66) <u>Kiwi</u>
8	(6.07) Harp	(7.49) <u>Pear</u>
9	(5.96) Clarinet	(7.41) <u>Chestnut</u>
10	(5.66) Organ	(6.42) <u>Pomegranate</u>

Conclusion

In this paper we have presented ARCLASS, a generalization of the Tversky's Ratio model of similarity which is grounded on:

1. an abstract representation of non-observable semantic features (the abstract relevance classes)
2. a reduction technique (NMF) which has been shown to be useful in modeling hidden semantic structures.

We have applied ARCLASS to the field of concept representations by analyzing a high-dimensional database of 254 (concepts) $\times$ 2619 (semantic features). ARCLASS extracted solutions which intuitively correspond to what humans call categories such as Animals, Vegetables, Furniture etc. What is more, ARCLASS results can be applied to explain some puzzling phenomena observed in the field of semantic memory disorders where it has been shown that some neurological patients may be impaired in one category (e.g. Living objects) but not in others (e.g. Non-living objects). In sum, ARCLASS may be considered as a new valuable tool to measure similarity patterns within featural frameworks.

Acknowledgments

The authors gratefully acknowledge Luigi Burigana for a critical reading of the manuscript and helpful suggestions.

References

Ashcraft, M. H. (1978). Property norms for typical and atypical items from 17 categories: A description and discussion. *Memory and Cognition, 6*, 227-232.

Benjamini, Y., & Hochberg, Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistic Society B*, 57, 298-300.

Capitani, E., Laiacona, M., Mahon, B., & Caramazza, A. (2003). What are the facts of semantic category-specific deficits? A critical review of the clinical evidence. *Cognitive Neuropsychology, 20*, 213-261.

Dell'Acqua, R., Lotto, L., & Job, R. (2000). Naming time and standardized norms for the Italian PD/DPSS set of 266 pictures: Direct comparisons with American, English, French, and Spanish published databases. *Behaviour Research Methods, Instruments and Computers*, *32*, 588-612.

Dixon, M.J., Piskopos, K., & Schweizer, T.A. (2001). Musical instruments naming impairments: The crucial exception to the living/non-living dichotomy in category specific agnosia. *Brain & Cognition, 43*, 158-164.

Elman, J.L. (1993). Learning and development in neural networks: the importance of starting small. *Cognition, 48*, 71-79.

Fishburn, P.C. (1970). Intransitive indifference in preference theory: a survey. *Operations Research, 18*, 207-228.

Goldstone, R.L., & Medin, D.L. (1994). The time course of comparison. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20*, 29-50.

Halford, G.S., Baker, R., McCredden, J.E., & Bain, J,D, (2005) How many variables can humans process? *Psychological Science, 16*, 70-76.

Humphreys, G.W., & Forde, E.M.E. (2001). Hierarchies, similarity and interactivity in object recognition: "Category-specific" neuropsychological deficits. *Behavioral and Brain Sciences, 24*, 453-509.

Kruschke, J.K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review, 99*, 22-44.

Lee, D., & Seung, H. (1999). Learning the parts of objects by non-negative matrix factorization. *Nature, 401*, 788-791.

Lee, D., & Seung, H. (2001). Algorithms for non-negative matrix factorization. *Advances in Neural Information Processing Systems, 13*, 556-562.

Love, B.C., Medin, D.L., & Gureckis, T.M. (2004). SUSTAIN: A network model of category learning. *Psychological Review, 111*, 309-332.

Marques, J. F. (2005). Naming from definition: The role of feature type and feature distinctiveness. *The Quarterly Journal of Experimental Psychology: Section A*, *58*, 603-611.

Mechelli A., Sartori, G., Orlandi, P., Price, C.J. (2005). Semantic relevance explains category effects in medial fusiform gyri. *Neuroimage*, in press.

Medin, D.L., Golstone, R.L., & Gentner, D. (1993). Respects for similarity. *Psychological Review, 100*, 254-278.

Miller, G.A. (1956). The magical number seven, plus or minus two: some limits on our capacity for processing information. *Psychological Review, 63*, 81-97.

Navarro, D.J., & Lee, M.D. (2004). Common and distinctive features in stimulus similarity: A modified version of the contrast model. *Psychonomic Bullettin & Review, 11*, 961-974.

Nosofsky, R.M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General, 115*, 39-57.

Salton, G. (1989). *Automatic text processing: The transformation, analysis, and retrieval of information by computer.* Reading, MA: Addison-Wesley.

Sartori, G., & Lombardi, L. (2004). Semantic relevance and semantic disorders. *Journal of Cognitive Neuroscience*, *16*, 439-452.

Sartori, G., Lombardi, L., & Mattiuzzi, L. (2005). Semantic relevance best predicts normal and abnormal name retrieval. *Neuropsychologia, 43*, 754-770.

Sartori G., Polezzi, D., Mameli, F., Lombardi L. (2005). Feature type effects in semantic memory: An event related potentials study. *Neuroscience Letters, 390*, 139-144.

Shepard, R.N. (1987). Toward a universal law of generalization for psychological science. *Science*, *237*, 1317-1323.

Siegel, S., Castellan, N.J. (1988). *Nonparametric statistics for the behavioral sciences* (2nd ed.). McGraw-Hill, New York.

Tenenbaum, J.B., & Griffiths, T.L. (2001). Generalization, similarity, and Bayesian inference. *Behavioral and Brain Sciences, 24*, 629-640.

Tversky, A. (1977). Features of similarity. *Psychological Review*, *34*, 327-352.