Acta Psychologica 133 (2010) 296–304

Contents lists available at ScienceDirect

Acta Psychologica

journal homepage: www.elsevier.com/locate/actpsy

Detecting intra- and inter-categorical structure in semantic concepts using HICLAS

Eva Ceulemans[a,*,1], Gert Storms[b,1]

[a] Department of Educational Sciences, University of Leuven, Vesaliusstraat 2, Box 3762, B-3000 Leuven, Belgium

[b] Department of Psychology, University of Leuven, Tiensestraat 102, Box 3721, B-3000 Leuven, Belgium

ARTICLE INFO

Article history:
Received 31 December 2008
Received in revised form 18 November 2009
Accepted 23 November 2009
Available online 30 December 2009

PsycINFO classification:
2240

Keywords:
Modeling
Semantic concepts
Inter-categorical structure
Intra-categorical structure

ABSTRACT

In this paper, we investigate the hypothesis that people use feature correlations to detect inter- and intra-categorical structure. More specifically, we study whether it is plausible that people strategically look for a particular type of feature co-occurrence that can be represented in terms of rectangular patterns of 1s and 0s in a binary feature by exemplar matrix. Analyzing data from the Animal and Artifact domains, we show that the HICLAS model, which looks for such rectangular structure and which therefore models a cognitive capacity of detecting feature co-occurence in large data bases of features characterizing exemplars, succeeds rather well in predicting inter- and intra-categorical structure.

1. Introduction

Ever since Rosch and her collaborators (Rosch & Mervis, 1975; Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976) published their influential studies on semantic concepts in the mid seventies, both the structure between categories and the structure within categories have been studied extensively in cognitive psychology (see Medin, Lynch, & Solomon, 2000; Murphy, 2002, for overviews). Moreover, an important topic of debate – on which we focus in this paper – is how people detect this inter- and intra-categorical structure.

1.1. Inter- and intra-categorical structure of semantic categories

Regarding inter-categorical structure, or the structure between categories, Rosch et al. (1976) focused on the hierarchical nature of many semantic concepts. The Animal category, for instance, falls apart into *mammals*, *birds*, *fish*, *insects*, *reptiles*, etc., that is, into mutually exclusive categories that are defined at a lower level of abstraction. Each of these lower level categories is further subdivided into less abstract categories. The *mammal* category, for instance, consists of *dogs*, *cats*, *cows*, *horses*, *elephants*, and so on, and categories such as *dogs* can even be further subdivided into *German shepards*, *poodles*, *rottweilers*, etc. Artifacts are likewise structured hierarchically. A particular object can be referred to as a jazz guitar, a guitar, a string instrument, a musical instrument, or an artifact.

Regarding intra-categorical structure, or the structure within categories (Rosch et al., 1976; Rosch & Mervis, 1975) convincingly showed that not all exemplars are equally good examples of a category. People agree rather well on how typical exemplars of common semantic categories are. For example, most people agree that a *cow* is a better instance of a *mammal* than a *whale* (even though they will agree that both of these Animal kinds are mammals) and that a *piano* is a better instance of a *musical instrument* than an *Indonesian gamelan* (McCloskey & Glucksberg, 1978). Moreover, ratings of typicality are quite consistent over time (Barsalou, 1987; Hampton, 2006). The graded structure of semantic categories, as measured by rated typicality, has been shown to predict performance in many other category-related tasks, such as inductive reasoning (Osherson, Smith, Wilkie, López, & Shafir, 1990), exemplar production (Storms, 2001), category naming (Storms, De Boeck, & Ruts, 2000), priming effects (Rosch, 1977), memory interference effects (Keller & Kellas, 1978), and response times from a speeded categorization task (Hampton, 1979). It thus seems that semantic categories display a stable internal structure that needs to be accounted for by any theory of concept representation.

Both the inter- and the intra-categorical structure are reflected in two aspects of semantic concepts: the extension and the intension. The extension of a concept corresponds to the class of entities that the concept refers to. The intension of a concept is the idea

* Corresponding author. Tel.: +32 16 325 881; fax: +32 16 325 934.
E-mail addresses: eva.ceulemans@ped.kuleuven.be (E. Ceulemans), gert.storms@psy.kuleuven.be (G. Storms).

[1] Eva Ceulemans and Gert Storms contributed equally to this paper.

0001-6918/$ - see front matter © 2009 Elsevier B.V. All rights reserved.
doi:10.1016/j.actpsy.2009.11.011

associated with the concept, or, the set of salient or psychologically important features that delineate the concept.

The inter-categorical structure is reflected extensionally in which entities belong to which of a series of (contrasting) categories. For instance, carrots and tomatoes belong to the *vegetable* category, while apples and oranges belong to the category of *fruits*. The intra-categorical structure is reflected extensionally in the extent to which a particular exemplar of a category is typical of the category. Pears are generally considered to be more typical for the *fruit* category than water melons. But one can also look at inter- and intra-categorical structure from an intensional point of view. The features belonging to the intension of a semantic concept X but not to the intension of another concept Y and the features belonging to the intension of Y but not to that of X are thus crucial to the differentiation between both concepts, or, to use different words, to the inter-categorical structure of both concepts. Tasting sweet, for instance, is important in distinguishing *fruits* from *vegetables*. Finally, the intra-categorical structure is reflected intensionally in the features that contribute in making exemplars more or less typical of a category. Though strawberries and blackberries are fruits that do not grow on trees, more typical fruit exemplars like apples and oranges do grow on trees.

Since semantic concepts are, as mentioned above, usually hierarchically structured, intra-categorical structure (i.e., which features apply to which exemplars) and inter-categorical structure (i.e., which features differentiate exemplars from non-exemplars) cannot be separated. Of course, if a category X, defined at a hierarchically higher level, falls apart into categories A, B, and C at a lower level, then intra-categorical structure of X reflects inter-categorical structure of A, B, and C.

1.2. Studying the cognitive capacity to detect inter- and intra-categorical structure

The question how people learn to assign stimuli to categories has drawn a lot of attention, which resulted in the development of a large number of (formal) categorization models that have been tested extensively in experiments with artificial categories. For an overview, see for instance, Ashby and Maddox (2005) and Smith and Minda (2000).

The related question of how people carve up common entities in the real world into more or less homogeneous categories is more difficult to answer. Many categories like Animal sorts and Artifact categories are learned by children in non-verbal ways at very early ages, which makes the process difficult to study in rigorously controlled settings. As a result, researchers are forced to study young children using indirect measures (Mandler, 2000), or take refuge in studying categorization of novel stimuli into well-known categories in adults (e.g., Ameel, Storms, Malt, & Sloman, 2005; Smits, Storms, Rosseel, & De Boeck, 2002; Storms, De Boeck, & Ruts, 2001).

Another way to study how the virtually infinite diversity of entities in the world gets carved up into categories is by developing models that mimic human categorization. Recently, Rogers and McClelland (2004) proposed a parallel distributed-processing approach to model the category learning process faced by children. In this paper, we take another approach and explore the possibility that people deduce the inter- and intra-categorical structure of semantic concepts from the correlational structure of psychologically salient features in the entities in the world (Storms & De Boeck, 1997; Storms, Van Mechelen, & De Boeck, 1994) and we model this process by using hierarchical classes analysis (HICLAS: De Boeck & Rosenberg, 1988).

Thus, the central idea to our approach is that the world outside us is structured in such a way that the entities are carved up into categories that are characterized by similar patterns of psychologically important features and, as Rosch and Mervis (1975) argued in the seventies, that intra-categorical structure arises from correlational structure *within* categories. Moreover, if people are sensitive to this correlational structure, they may learn to use some features to divide the entities in the world into separate categories, while using other features to determine typicality within the categories and ignoring still other features when structuring categories at a particular level of abstraction (note, for instance, that the model of Zeigenfuse and Lee (2010) incorporates the assumption that people ignore some features).

The way we approach the problem of predicting inter and intra categorical structure is by analyzing exemplar by feature data and by looking for dense regions in the data, after reorganizing rows and columns of the exemplar by feature matrices. Such dense regions reflect (near) monothetic or rectangular patterns and, as will be explained below, they can be found by applying the disjunctive HICLAS model (De Boeck & Rosenberg, 1988). As such, the HICLAS model mimics the potential cognitive ability of finding correlational structure in the data.

In the remainder of this paper, we will first briefly describe the HICLAS model and then continue with a presentation of the results of the HICLAS analyses of two data sets: one consisting of exemplar by feature data from the Animal domain and another data set with Artifact data. The paper will be concluded with reflections on the cognitive implications of the use of data analytic tools such as HICLAS to study inter- and intra-categorical structure.

2. HICLAS

The HICLAS model (De Boeck & Rosenberg, 1988) is a structural model for a binary I exemplars by J features data matrix $\mathbf{D}$. In the following paragraphs, we will use the hypothetical four exemplars by six features matrix in Table 1 as a guiding example. In this table, a value of 1 indicates that the corresponding exemplar is characterized by the feature in question; a value of 0 implies that this is not the case. For instance, it can be read that a zebra is warm-blooded, but has no air sacs.

Given a binary I exemplars by J features data matrix and a rank R, HICLAS approximates the data by a binary I exemplars by J features model matrix $\mathbf{M}$, such that the following loss function is minimized:

$$L = \sum_{i=1}^{I} \sum_{j=1}^{J} (d_{ij} - m_{ij})^2, \tag{1.1}$$

subject to the restriction that, permuting the exemplars and features, $\mathbf{M}$ contains R possibly overlapping rectangles of ones. This approximation is achieved by reducing the exemplars and features to R overlapping clusters, called bundles, using an alternating least squares or simulated annealing procedure (Ceulemans et al., 2007). From the exemplar and feature bundles, the exemplar by feature model matrix can be computed by means of the following association rule:

$$m_{ij} = \bigoplus_{r=1}^{R} e_{ir} f_{jr}, \tag{1.2}$$

Table 1
Hypothetical exemplar by feature data matrix **D**.

Exemplar	Feature					
	Has nipples	*Breastfeeds*	*Warm-blooded*	*Soft*	*Feathers*	*Air sacs*
Whale	0	1	1	0	0	0
Zebra	1	1	1	1	0	0
Blackbird	0	0	1	1	1	1
Woodpecker	0	0	1	1	1	1

Table 2
HICLAS model matrix **M** of rank 2 for the data in Table 1.

Exemplar	Feature					
	Has nipples	*Breastfeeds*	*Warm-blooded*	*Soft*	*Feathers*	*Air sacs*
Whale	1	1	1	1	0	0
Zebra	1	1	1	1	0	0
Blackbird	0	0	1	1	1	1
Woodpecker	0	0	1	1	1	1

where $\oplus$ denotes the Boolean sum (i.e., $1 \oplus 1 = 1$), and e_{ir} and f_{jr} indicate the entries of the exemplar bundle matrix **E** and the feature bundle matrix **F**, respectively. For instance, Table 2 shows a HICLAS model matrix of rank 2 for the data in Table 1, which is obtained by reducing the exemplars and features to two bundles. The corresponding exemplar and feature bundle matrices **E** and **F** are given in Table 3. Comparing Tables 1 and 2, one can derive that this HICLAS model has a loss function value of 2.

The bundles of the exemplars and the features are linked to one another, implying that exemplars that belong to a specific exemplar bundle are characterized by all features that belong to the corresponding feature bundle. For instance, Table 2 shows that in the HICLAS model *blackbird* and *woodpecker* are characterized by *is warm-blooded*, *is soft*, *has feathers*, and *has air sacs*. This is reflected in Table 3, as *blackbird* and *woodpecker* belong to the second exemplar bundle and *is warm-blooded*, *is soft*, *has feathers*, and *has airs sacs* constitute the second feature bundle.

The one-to-one correspondence of the exemplar and feature bundles imply that exemplars and features that belong to linked exemplar and feature bundles, constitute a rectangle of 1s in the data matrix. Similarly, exemplars and features that do not belong to the corresponding exemplar and feature bundles, form a rectangle of 0s in the data matrix. For instance, in Table 3, *blackbird* and *woodpecker* belong to the second exemplar bundle only, whereas *has nipples* and *breastfeeds* belong to the first feature bundle only. This is reflected in Table 2 in that these two exemplars and two features constitute a rectangle of 0s. It can be concluded that whereas each rectangle of 1s represents a category, the corresponding exemplar bundle represents the extension of the category and the corresponding feature bundle represents its intension. For instance, in Table 3, the first exemplar and feature bundles constitute the category of mammals; *zebra* and *whale* are the extension of this category, and *has nipples*, *breastfeeds*, *is warm-blooded*, and *is soft* are the intension.

2.1. Hierarchically organized classifications

An important feature of the HICLAS model is that it includes hierarchically organized classifications of the exemplars and features. This hierarchical structure is obtained by applying a closure operation to the bundle matrices **E** and **F** at the end of the alternating least squares or simulated annealing procedure (see e.g., Ceulemans et al., 2007); this closure operation does not alter the loss function value of the obtained solution. With respect to the classification, exemplars that are characterized by the same set of features constitute an exemplar class, whereas features that apply to the same set of exemplars constitute a feature class. These classifications of the exemplars and the features are represented in the bundles, in that all the elements of a class belong to the same set of bundles. For instance, from Table 2 it can be read that *zebra* and *whale* are characterized by the same set of features (*has nipples*, *breastfeeds*, *is warm-blooded*, *is soft*); hence, these two animals belong to the same exemplar class. This is reflected in Table 3 in that both *zebra* and *whale* belong to the first bundle, but not to the second one. Similarly, *has feathers* and *has air sacs* constitute a feature class, as both features apply to the same two exemplars (i.e., *blackbird*, *woodpecker*). As such, *has feathers* and *has air sacs* have identical bundle patterns in Table 3.

With respect to the hierarchical organization of the classifications of the exemplars and the features, an exemplar class is hierarchically lower than another exemplar class if the set of features that apply to the first class is a subset of those that apply to the latter class. Similarly, a feature class is hierarchically lower than another feature class if the set of exemplars that are characterized by the first class is a subset of the set of exemplars that are characterized by the latter class. The hierarchical relations between the exemplar classes and the feature classes are also represented in the bundles, in that the bundle pattern of a hierarchically lower class is a subset of that of a hierarchically higher class. For instance, Table 2 shows that *has feathers* and *has air sacs* are hierarchically lower than *is warm-blooded* and *is soft*, because the first two features only apply to *blackbird* and *woodpecker*, whereas the last two features apply to all exemplars. This hierarchical relation is represented in Table 3 in that *has feathers* and *has air sacs* only belong to the second bundle, whereas *warm-blooded* and *is soft* belong to both bundles.

Table 3
Bundle matrices **E** and **F** of rank 2 for the data in Table 1.

Exemplar	Bundles		Feature	Bundles	
	I	II		I	II
Whale	1	0	*Has nipples*	1	0
Zebra	1	0	*Breastfeeds*	1	0
Blackbird	0	1	*Is warm-blooded*	1	1
Woodpecker	0	1	*Is soft*	1	1
			Has feathers	0	1
			Has air sacs	0	1

2.2. Graphical representation

A graphical representation of the HICLAS model in Tables 2 and 3 is given in Fig. 1. In this figure, the hierarchical classifications of the exemplars and the features are drawn in the upper and lower half of the representation, respectively, where the feature hierarchy is represented upside down. The classes are indicated by the boxes and the hierarchical relations between the classes by the lines between the boxes. Finally, the zigzags represent the linking structure between the exemplar bundles and the feature bundles.

2.3. Model selection

In practice, the structure of binary exemplar by feature data can only be perfectly reconstructed (i.e., loss function value equals zero) by means of a HICLAS model with a large number of bundles. Such models are too complex to be useful, however. Therefore, one will usually look for a model that describes the data well without

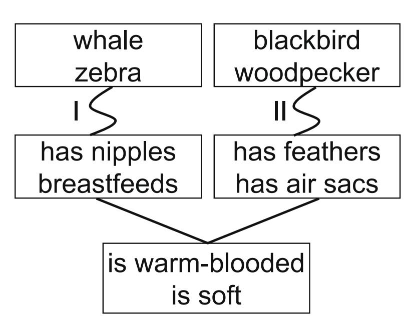

Fig. 1. Graphical representation of the HICLAS model in Table 2.

being overly complex. This is achieved by fitting HICLAS models with increasing numbers of bundles R to the data. Subsequently, the solution that best balances fit to the data (i.e., low L-value) and complexity (i.e., number of bundles R) is selected by applying some model selection procedure.

Recently, a range of such procedures has been proposed, including a numerical convex hull-based extension of the well-known scree test (Ceulemans & Van Mechelen, 2005), a pseudo-binomial test (Leenen & Van Mechelen, 2001), and, recasting HICLAS in probabilistic terms, a pseudo-AIC criterion (Ceulemans & Van Mechelen, 2005), Bayes factors and posterior predictive checks (Leenen, Van Mechelen, Gelman, & De Knop, 2008). In this paper, we only consider the original, deterministic HICLAS model, and will therefore use the numerical convex hull-based procedure and the pseudo-binomial test.

The numerical convex hull-based procedure selects the solution on the lower boundary of the convex hull of a number of bundles vs. loss function value plot, after which the decrease in loss function levels off (Ceulemans & Kiers, 2006; Ceulemans & Van Mechelen, 2005). The pseudo-binomial test selects the smallest rank or amount of bundles R, for which the probability of observing a value X smaller than the loss function value of the model of rank $R + 1$ is greater than some small number, say, .0001 given that X is binomially distributed as follows: $X \sim \text{bin}(IJ, L_R/IJ)$, with IJ equaling the number of cells in the data matrix and L_R indicating the loss function value of the solution with R bundles (Leenen & Van Mechelen, 2001).

2.4. Fit measures

Apart from the loss function value, which boils down to the number of discrepancies between the data and the model (i.e., the number of cells with a 0 in the data and a 1 in the model or vice versa), another fit measure that is often used in HICLAS analysis is the Jaccard goodness-of-fit index (Sneath & Sokal, 1973), which takes values between 0 and 1:

$$J = \frac{n_{D=1,M=1}}{n_{D=1,M=1} + n_{D=0,M=1} + n_{D=1,M=0}}, \tag{1.3}$$

where $n_{D=1,M=1}$ indicates the number of cells with a 1 in the data and a 1 in the model, $n_{D=0,M=1}$ indicates the number of cells with a 0 in the data and a 1 in the model, and $n_{D=1,M=0}$ indicates the number of cells with a 1 in the data and a 0 in the model. This Jaccard index can be calculated for the overall model as well as for each exemplar and each feature separately. Through the varying Jaccard values for the exemplars and the features, the HICLAS model leaves room for gradedness, both on the extensional and on the intensional side (Storms et al., 1994). Specifically, the Jaccard indices for the exemplars can be considered a measure of prototypicality and the Jaccard indices for the features a measure of category relevance. For instance, the Jaccard indices for *zebra* and *whale* equal $4/(4 + 0 + 0) = 1$ and $2/(2 + 2 + 0) = .5$ respectively, indicating that *zebra* is a more prototypical mammal than *whale*. Note that the Jaccard index is closely related to the similarity index s_{ij} between a pair of objects i and j in the ratio model of Tversky (1977), which takes the common and distinctive features of both objects into account.

2.5. Relations to other models

In the past decades, many models and associated algorithms have been proposed that can be applied to binary exemplar by feature data in order to obtain a clustering of the exemplars as well as the features (for an overview of two-mode clustering methods, see Van Mechelen, Bock, & De Boeck, 2004). The models differ in many respects, however: Are the models deterministic or stochastic? Does the induced clustering of the exemplars and the features take the form of a partition or is it an overlapping clustering? Is the number of exemplar and feature clusters fixed throughout the estimation procedure or is it dynamically updated?

Comparing HICLAS to two two-mode clustering techniques that have recently been proposed in the categories and concepts literature, i.e., the Infinite Relational Model (Kemp et al., 2006) and the CrossCat model (Shafto et al., 2006) shows that HICLAS differs in all three respects from these models: whereas the Infinite Relational Model and the CrossCat model are stochastic, the original HICLAS model that is used in this paper, is deterministic (for a minimal stochastic extension of HICLAS, see Leenen et al., 2008). HICLAS yields an overlapping clustering of the exemplars and the features, which also implies a partitioning of both sets however, based on the subset of the overlapping clusters to which an exemplar or feature belongs; applying the Infinite Relational Model and the CrossCat model results in partitions of exemplars and features. Finally, the existing HICLAS algorithms require the specification of the number of overlapping clusters (bundles), whereas the Infinite Relational model and the CrossCat model determine the number of clusters dynamically.

3. Applying HICLAS to detect inter- and intra-categorical structure in semantic concepts

Two data sets out of the Leuven Natural Concept Database, taken from De Deyne et al. (2008) were analyzed with the disjunctive HICLAS model (Ceulemans, Van Mechelen, & Leenen, 2007; De Boeck & Rosenberg, 1988). The first data set consisted of a 129 exemplar by 225 feature matrix for the Animal domain (i.e., the Type IV data set from De Deyne et al. (2008)). The exemplar set contained a sample of 30 *mammals*, 30 *birds*, 23 *fish*, 26 *insects*, and 20 *reptiles* and *amphibians*, representative in terms of presumed typicality, ranging from very atypical exemplars to very typical exemplars of these categories, but with the restriction that all the selected exemplars had to be familiar to the vast majority of an adult population in the Dutch-speaking part of Belgium (note that the participants in the generation study of Storms, 2001, on which the exemplar selection was based, did not know the difference between reptiles and amphibians very well, generating a lot of overlapping exemplars for both categories. Therefore, these two categories were treated as a single category).

The features were selected from a feature generation task in which participants were asked to name 10 features that define (technically or loosely speaking) each of the five Animal categories. They were encouraged to include different kinds of features, including physical, perceptual, and functional features as well as encyclopaedic knowledge. For each of the Animal categories, 20 different participants generated features. For the categories *birds*, *fish*, *insects*, *mammals*, and *reptiles*, respectively 52, 58, 70, 60, and 70 features were selected for inclusion in the exemplar by feature matrix, but note that some features overlapped.

The second data set consisted of a 166 exemplar by 300 feature matrix for the domain of Artifacts (i.e., again the Type IV data set from De Deyne et al. (2008)). Unlike the Animal categories, the Artifact categories are not mutually exclusive and no generally accepted delineation of these categories exists (unlike for the Animal categories, for which it is assumed that there is a biological taxonomy of non-overlapping categories). Several objects were generated frequently by participants as instances of more than 1 of the 6 Artifact categories under study. Assigning exemplars to the categories where they were generated with the highest frequencies, the exemplar set included a sample of 27 *musical instruments*, 27 *tools*, 30 *vehicles*, 29 *clothing items*, 33 *kitchen utensils*, and 20 *weapons*. Again these exemplars were selected to be representative in terms of presumed typicality, ranging from very atypical exemplars to very typical exemplars of these categories, but with the restriction that all of the selected exemplars had to be familiar to the vast majority of an adult population in the Dutch-speaking part of Belgium.

The features were again selected from a feature generation task in which participants were asked to name 10 features that define (technically or loosely speaking) each of the six Artifact categories. As for the Animal data set, they were encouraged to include different kinds of features, including physical, perceptual, and functional features as well as encyclopaedic knowledge. For the categories *clothing*, *kitchen utensils*, *musical instruments*, *tools*, *vehicles*, and *weapons*, respectively 71, 73, 74, 79, 62, and 54 features were selected for inclusion in the exemplar by feature matrix, but note that some features overlapped.

The entries in the analyzed Animal and Artifact exemplar by feature matrices were the number of subjects (out of four) who judged (for every feature-exemplar pair) whether the feature characterizes the exemplar or not. The reliability of the two exemplar by feature matrices was estimated by De Deyne et al. (2008) using the Spearman–Brown split-half technique. As there are only three different ways to divide four subjects in two groups of two, all three possible splits were evaluated, resulting in reliability estimates of .89, .90, and .87 for the Animal matrix, and .85, .84, and.84 for the Artifact matrix.

As a measure of the intra-categorical structure, we used typicality ratings in the eleven studied categories, also gathered by De Deyne et al. (2008). The reliability of these ratings was quite high, with estimated values all above .90, except for the category of *insects*, where the estimated value was .87.

3.1. Results and discussion of the analysis of the animal data set

The 129 animals by 225 features data matrix was dichotomized using a majority rule, that is, by replacing values ranging from 0 to 1 by 0 and values ranging from 2 to 4 by 1. HICLAS solutions including 1–8 bundles were obtained, of which the loss function values and the Jaccard goodness-of-fit indices are shown in Fig. 2. Both the numerical convex hull-based model selection procedure (Ceulemans & Kiers, 2006; Ceulemans & Van Mechelen, 2005) and the pseudo-binomial test (Leenen & Van Mechelen, 2001) indicated the selection of the HICLAS solution with five bundles. Fig. 3 shows a graphical representation of this solution. In this figure, the Animal classes are labeled by indicating to which of the

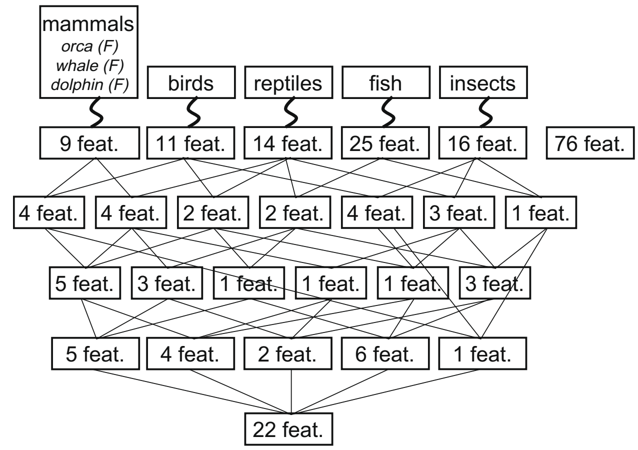

Fig. 3. HICLAS model with five bundles for the Animal data set.

five underlying categories – *mammals*, *birds*, *reptiles*, *fish*, *insects* – (most of) the Animals in the class belong. Additionally, Animals for which the class and category membership do not correspond are printed in italics, with the category membership indicated between brackets.

Inspecting Fig. 3, one immediately notices that the exemplar structure, consisting of five classes only, is much simpler than the feature structure, which is composed of 25 classes. Moreover, the five Animal classes, which each belong to only one bundle, perfectly correspond to the underlying categories – *mammals*, *birds*, *reptiles*, *fish*, and *insects* – with three exceptions only: *orca*, *whale*, and *dolphin* are assigned to the mammal bundle instead of the fish bundle. This makes perfect sense as, although these three animals were generated as exemplars of fish in the exemplar generation task (De Deyne et al., 2008), according to a biological taxonomy they are mammals and not fish. It can be concluded that HICLAS yields a perfect delineation of the Animal categories under study and as such succeeds in recovering the inter-categorical structure.

Though we concentrate on the exemplar structure in this paper, it is important to notice that the feature structure is complicated, but clearly interpretable. Examples of good-fitting features (with a Jaccard index above .90) in the bundle specific feature classes corresponding to *mammals*, *birds*, *reptiles*, *fish*, and *insects*, were *has nipples*, *has feathers*, *is cold-blooded*, *breathes through gills*, and *has feelers*, respectively. The features in the hierarchically higher

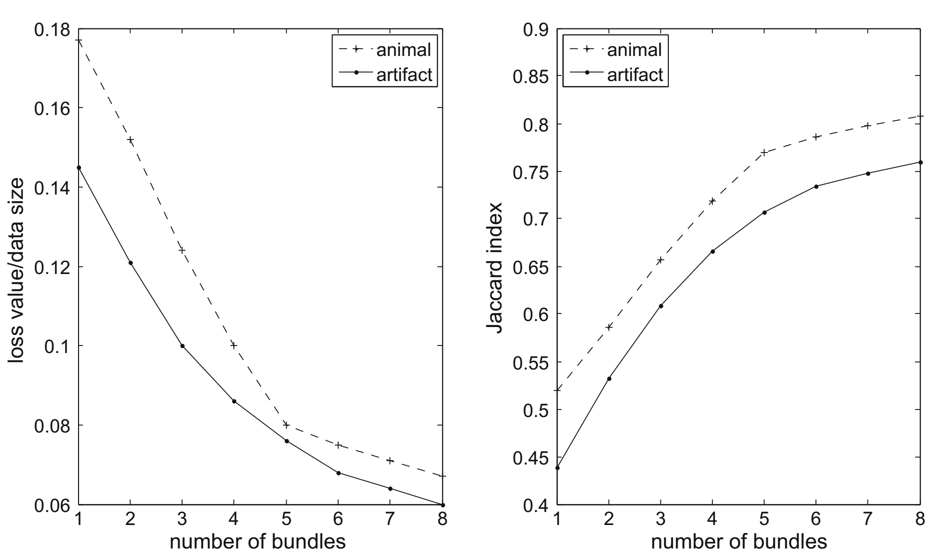

Fig. 2. Loss function values (divided by data size) and Jaccard indices of the HICLAS solutions with 1 up to 8 bundles, for the Artifact and the Animal data.

classes were equally well-interpretable. *Lays eggs*, for instance, shows up in a feature class that spans all categories except *mammals*, and features like *eats* and *has brains* are assigned to the feature class that applies to all animals.

As stated in the introduction, not all exemplars are equally good examples of a category, implying that semantic categories have a graded structure. This intra-categorical structure is reflected in the rated typicality of the exemplars for the categories. To study whether HICLAS retrieves this intra-categorical structure, the Jaccard goodness-of-fit indices of the animals within a category were correlated with the corresponding typicality ratings. Note that *orca*, *whale* and *dolphin* were excluded from these calculations, because their typicality ratings concerned typicality for the category *fish* instead of typicality for the category *mammal*. For *reptiles* ($r = .51$, $p = .02$), *insects* ($r = .43$, $p = .03$), *fish* ($r = .57$, $p = .01$), and *birds* ($r = .66$, $p < .0001$), the obtained correlations were significantly positive. For *mammals*, however, the correlation coefficient amounted to .04 ($p = .85$) only.

A possible explanation for this low correlation may be our familiarity with the mammal category. As Barsalou (1987) showed, typicality is highly correlated with familiarity as expressed in word frequency. In order to control for the effect of word frequency, we first predicted typicality in the *mammal* category from (the logarithm of) word frequency (De Deyne et al., 2008) and then correlated the residuals with the Jaccard goodness-of-fit indices from the HICLAS analysis. As a result, the predictive correlation rose to .36, a value that is nearly significant ($p = .06$).

Another consequence of our extensive familiarity with *mammals* is our detailed knowledge of this collection of species, which makes the *mammal* category less homogeneous than the other (relatively lesser known) Animal categories. As we will explain below (see Section 4), the resulting presence of subcategories in the *mammal* category may also account for the decreased correlation between fit and typicality.

3.2. Results and discussion of the analysis of the Artifact data set

Like the Animal data set, the 166 artifacts by 300 features data matrix were dichotomized by replacing values ranging from 0 to 1 by 0 and values ranging from 2 to 4 by 1. The dichotomized data matrix was analyzed with the HICLAS algorithm with the number of bundles varying from 1 to 8. Applying the numerical convex hull-based procedure and the pseudo-binomial test to Fig. 2, which shows the loss function values and the Jaccard goodness-of-fit indices of the eight resulting HICLAS solutions, resulted in the selection of the solution with six bundles. Fig. 4 shows a graphical representation of this solution; note that only the hierarchical classification of the artifacts is displayed, as the hierarchical classification of the 300 features is very complex (i.e., it consists of 33 classes).

With respect to inter-categorical structure, it can be read from Fig. 4 that some Artifact categories can be more easily delineated by means of HICLAS than others: Whereas *musical instruments* and *clothing* are clearly separated from the other categories, *vehicles*, *weapons*, *tools*, and *kitchen utensils* are slightly intertwined. This is no surprise as the Artifact categories under study are not mutually exclusive and no generally accepted delineation of these categories exists. Indeed, De Deyne et al. (2008) report that in the Artifact generation task multiple artifacts were generated frequently by participants as instances of more than one of the six Artifact categories. For instance, *axe* and *rope* were generated as instances of *weapons* as well as *tools*. Note also that, as in the analysis of the Animal data, the categories detected by HICLAS are, in terms of the vertical structure of semantic concepts (Rosch et al., 1976) defined at the superordinate level.

As in the analysis of the Animal data, the feature structure is complicated, but clearly interpretable. Examples of features with a Jaccard index above .90 in the feature classes that are linked to *kitchen utensils*, *musical instruments*, *tools*, and *clothing* only, were *is used by cooks*, *is used in orchestras*, *is used to work with*, and *is sewn*, respectively. The features in the hierarchically higher classes were again easily interpretable. *Is a good invention*, for instance, is classified in a feature class that spans all categories except *weapons*, and features like *can wear off* and *is available in different types* are assigned to the feature class that applies to all artifacts.

With respect to intra-categorical structure, correlation coefficients were computed between the typicality ratings of the artifacts and their Jaccard goodness-of-fit indices. In these calculations only the artifacts that belong to the hierarchically lowest Artifact classes (i.e., the classes with the category labels) were used, thus excluding artifacts that were assigned to more than one bundle and also discarding artifacts for which the typicality ratings did not concern the category to which they were assigned to by HICLAS (e.g., *lawnmower*, *cart*, *apron*). For *clothing* ($r = .64$,

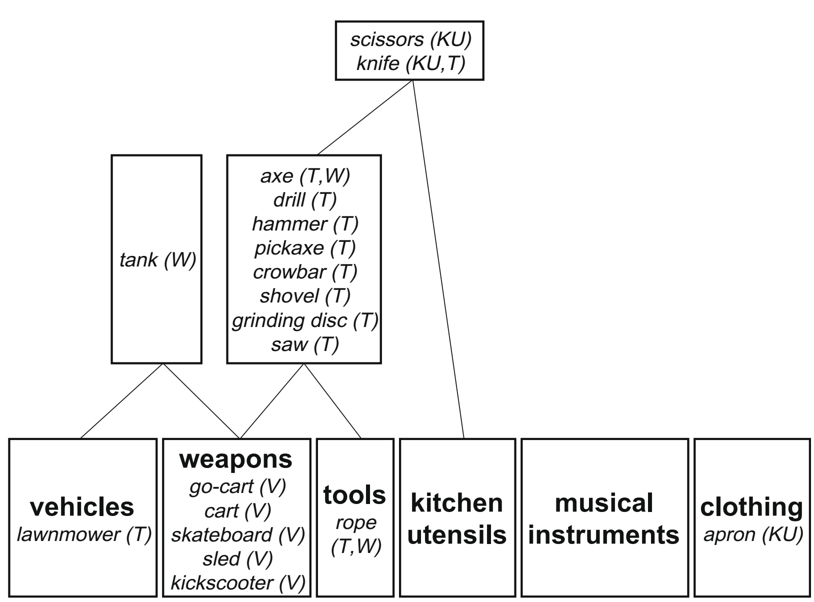

Fig. 4. HICLAS model with six bundles for the Artifact data set.

$p = .0002$) and *tools* ($r = .60$, $p = .0049$) the obtained correlations were significantly positive. For the other categories, however, the correlation coefficients were not significant, implying that the intra-categorical structure could not be predicted on the basis of the HICLAS solution.

4. General discussion

In this paper, we investigated the hypothesis that people use feature correlations to detect inter- and intra-categorical structure. More specifically, we studied whether it is plausible that people use a strategy of looking for a particular type of feature co-occurrence that is reflected by rectangular patterns in a binary feature by exemplar matrix. Analyzing data from the Animal and Artifact domains, we showed that the HICLAS model, which looks for such rectangular structure and which therefore models a cognitive capacity of detecting feature co-occurence in large data bases of features characterizing exemplars, succeeds rather well in predicting structure between categories, and even, to a certain degree, within categories.

Regarding inter-categorical structure, HICLAS was capable of splitting up a large number of animals in the correct number of underlying superordinate categories and of classifying all animals into the appropriate category. Interestingly, the model even succeeded in correctly assigning animals such as *orca*, *whale*, and *dolphin* to the *mammal* category, rather than to the *fish* category, in which these animals would be categorized if the judgments were solely based on superfluous perceptual and behavioral features. Furthermore, even though nothing prevented the model from classifying exemplars into hierarchically higher exemplar classes that correspond to the intersection of the base categories (i.e., the exemplar classes that belong to one bundle only), it assigned each of the 129 animals to one of the base categories, just like human subjects do when classifying living creatures in these categories. In line with common sense expectations, the feature structure consisted of clearly interpretable feature classes that apply exclusively to a single base category as well as feature classes that apply to multiple categories.

Looking at the results of the six-bundle solution for the artifacts, it is obvious that the base categories used to compose the stimulus set of the Artifact domain can be clearly distinguished. Unlike in the solution of the Animal domain, several exemplar stimuli were classified as belonging to multiple categories. The fact that HICLAS puts forward mutually exclusive categories in the Animal domain and overlapping categories in the Artifact domain fits nicely with findings from the literature that Artifact and natural kind categories have a different nature in this respect (Ruts, Storms, & Hampton, 2004; Sloman & Malt, 2003). The feature structure of the Artifact solution was again complicated, but clearly interpretable.

Regarding intra-categorical structure, the fit of the Animal exemplars within the base category to which they were assigned, succeeds in predicting typicality well for four of the five categories. Only in the *mammal* category the prediction turned out to be poor. We hypothesized that this lack of predictive power resulted from a relatively more elaborate knowledge, which results in a further detailed structure in this category. More specifically, we think that the fit of an exemplar within a bundle-specific class reflects typicality only if the corresponding category is rather homogeneous. If, on the contrary, this category falls apart into several subcategories, that is, if some of the relevant features for the category function to distinguish between the subcategories, then the applicability of the relevant feature set as a whole will poorly predict the typicality ratings of the exemplars for the category. Phrased yet differently, we hypothesize that typicality ratings in a heterogeneous category reflect a complicated relation towards one or more of the subcategories that does not straightforwardly correspond to the number of relevant category features that apply to the exemplar.

Looking in detail at the scatter plot of the typicality ratings versus the Jaccard goodness-of-fit indices of the exemplars, we saw that, for most categories with non-significant correlations, different subclusters could be clearly distinguished. For *kitchen utensils*, for instance, the group of electrical kitchen appliances yielded relatively low goodness-of-fit values in the HICLAS solution, but was rated as highly typical (see Fig. 5). The other exemplars showed a fairly linearly increasing pattern between fit and typicality that was broken by this subcategory of electric appliances. Likewise, for *weapons*, a fair correlation between fit and rated typicality was broken by a subcategory of fire arms with relatively low fit values, but that was rated as highly typical. A similar pattern was not immediately obvious in the *mammal* data, but because this Animal category is a very familiar one, about which we have detailed knowledge, it probably falls apart into a much larger number of subcategories.

In order to further test the hypothesis that the lack of correlation between typicality and goodness-of-fit is – at least partly – caused by the presence of clearly separate subcategories, we looked at the HICLAS solution with just one bundle for the Animal data. This means that the (very heterogeneous) Animal category, which – as we showed earlier – falls apart into several subcategories, and is treated as a single concept. Ratings of the typicality of all 129 animals for the category of Animals were gathered from 11 participants. They rated typicality on a 7-point scale, with a 1-value referring to very atypical and a 7-value referring to very typical animals. In line with our expectations, the correlation between rated typicality and goodness-of-fit within this heterogeneous Animal category were very low ($r = .04$).

4.1. Relating the HICLAS model to theories about semantic concept representation

HICLAS (Ceulemans, Van Mechelen, & Leenen, 2007; De Boeck & Rosenberg, 1988) is a data analytic technique for binary two-way two-mode data that can be used to obtain hierarchical classifications of the elements of both modes. One may ask the question why the results of a HICLAS analysis can be interpreted as evidence for a particular type of cognitive processing of semantic concepts. As stated before, we entertain the hypothesis that people use a particular type of feature correlations to detect inter- and intra-

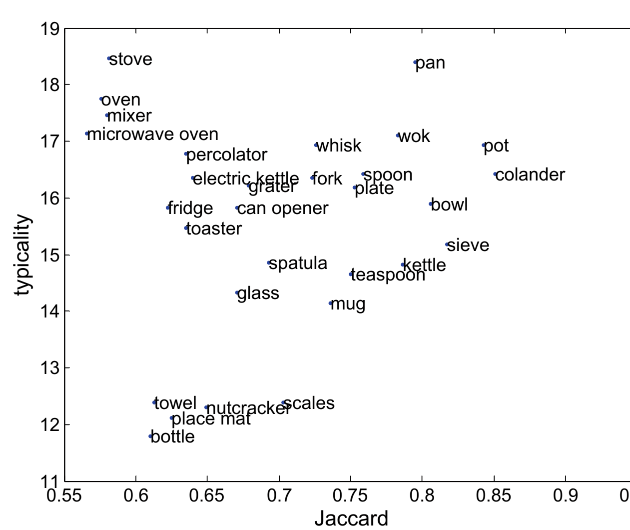

Fig. 5. Scatter plot of the typicality ratings versus the Jaccard goodness-of-fit indices of the kitchen utensils.

categorical structure. The type of feature correlations that HICLAS is sensitive to corresponds to a monothetic view on categories. In other words, if the rectangular patterns that HICLAS looks for are 'perfect' rectangles (i.e., rectangles constituted by a complete feature set that applies perfectly to a complete exemplar set), then the application of HICLAS corresponds to the classical view (Sutcliffe, 1993), where concepts are assumed to be represented by defining (i.e., singly necessary and jointly sufficient) features. However, as is generally known, this classical view on semantic concepts has been proven wrong in lots of empirical studies (Murphy, 2002; Smith & Medin, 1981). Completely in line with this, perfect rectangles will, of course, allow no differential predictions of typicality, as every exemplar within a category would show a perfect goodness-of-fit index.

Despite the overwhelming evidence against classically defined concepts, a weaker version of the same idea has been proposed by several authors, in which concepts are assumed to be represented by characteristic instead of defining features. Such characteristic features are not strictly necessary for every category exemplar, but they have substantial probabilities of occurring in the different instances of the concept. For examples of such a view, see for instance, Rosch and Mervis' (1975) family resemblance view and Hampton's (1979) prototype view. Also, when assuming that semantic concepts consist of relatively coherent clusters of exemplars (and equating the animal names with the exemplar level), the so-called exemplar view on semantic concepts (Nosofsky, 1986; Storms, 2004) is also compatible with approximate rectangles.

When the HICLAS model is used to parsimoniously describe the important patterns in binary data, the input exemplar by feature matrix is not perfectly, but approximately, represented by a HICLAS solution with a limited number of bundles. This implies that the rectangles that HICLAS looks for are not perfect rectangles, but approximate rectangles. Such approximate rectangles correspond to semantic concepts being represented by characteristic features. Furthermore, the goodness-of-fit of the exemplars (which is a function of the number of 1s in the discovered rectangles and of the 0s in the non-rectangular areas of the matrix) should then predict rated typicality within the concepts. As was shown in our study, the goodness-of-fit of the exemplars does correlate significantly to rated typicality and thus corresponds with the intra-categorical structure, at least in sufficiently homogeneous semantic concepts.

Finally, one may wonder whether our use of HICLAS as a cognitive model implies that we assume that concepts are organizational principles of an 'objective' material reality. In other words, do we adhere to a radical realist ontological view on concepts, in which the structure between and within concepts is given by the outside world (Van Mechelen, De Boeck, Theuns, & Degreef, 1993), rather than assuming that the concepts are mental representations, as is usually held by cognitive scientists? A fundamental answer to this question would require philosophical arguments that cannot be dealt with in this paper. Suffice it to say that the features on which the input data are based were taken from a feature generation study in which participants selected relevant properties of the studied categories, and that this selection process roots our usage of the HICLAS model at least partly at the mental level, rather than solely at the objective material reality.

Acknowledgements

The research reported in this paper was partly sponsored by Grants OT/05/27 of the Leuven Research Council and G.0513.08 of the 'Fonds voor Wetenschappelijk Onderzoek – Vlaanderen' to Gert Storms. The contributions to this special issue were first presented in a workshop, organized by the Research Group on Concepts and Categories of the Psychology Department, K.U. Leuven, in June 2008, which was sponsored by the 'Fonds voor Wetenschappelijk Onderzoek – Vlaanderen' (K.7.036.08N).

References

Ameel, E., Storms, G., Malt, B., & Sloman, S. A. (2005). How bilinguals solve the naming problem. *Journal of Memory and Language, 52*, 309–329.
Ashby, F. G., & Maddox, W. T. (2005). Human category learning. *Annual Review of Psychology, 56*, 149–178.
Barsalou, L. W. (1987). The instability of graded structure: Implications for the nature of concepts. In U. Neisser (Ed.), *Concepts and conceptual development: Ecological and intellectual factors in categorization*. Cambridge: Cambridge University Press.
Ceulemans, E., & Kiers, H. A. L. (2006). Selecting among three-mode principal component models of different types and complexities: A numerical convex hull based method. *British Journal of Mathematical & Statistical Psychology, 59*, 133–150.
Ceulemans, E., & Van Mechelen, I. (2005). Hierarchical classes models for three-way three-mode binary data: Interrelations and model selection. *Psychometrika, 70*, 461–480.
Ceulemans, E., Van Mechelen, I., & Leenen, I. (2007). The local minima problem in hierarchical classes analysis: An evaluation of a simulated annealing algorithm and various multistart procedures. *Psychometrika, 72*, 377–391.
De Boeck, P., & Rosenberg, S. (1988). Hierarchical classes: Model and data analysis. *Psychometrika, 53*, 361–381.
De Deyne, S., Verheyen, S., Ameel, E., Vanpaemel, W., Dry, M., Voorspoels, W., et al. (2008). Exemplar by feature applicability matrices and other Dutch normative data for semantic concepts. *Behavioral Research Methods, 40*, 1030–1048.
Hampton, J. A. (1979). Polymorphous concepts in semantic memory. *Journal of Verbal Learning and Verbal Behavior, 18*, 441–461.
Hampton, J. A. (2006). Concepts as prototypes. In B. H. Ross (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 46, pp. 79–113).
Keller, D., & Kellas, G. (1978). Typicality as a dimension of encoding. *Journal of Experimental Psychology: Human Learning & Memory, 4*, 78–85.
Kemp, C., Tenenbaum, J. B., Griffiths, T. L., Yamada, T., & Ueda, N. (2006). Learning systems of concepts with an infinite relational model. In *Proceedings of the 21st national conference on artificial intelligence.*
Leenen, I., & Van Mechelen, I. (2001). An evaluation of two algorithms for hierarchical classes analysis. *Journal of Classification, 18*, 57–80.
Leenen, I., Van Mechelen, I., Gelman, A., & De Knop, S. (2008). Bayesian hierarchical classes analysis. *Psychometrika, 73*, 39–64.
Mandler, J. M. (2000). Perceptual and conceptual processes in infancy. *Journal of Cognition and Development, 1*, 3–36.
McCloskey, M., & Glucksberg, S. (1978). Natural categories: Well-defined or fuzzy sets? *Memory and Cognition, 6*, 462–472.
Medin, D. L., Lynch, E. B., & Solomon, K. O. (2000). Are there kinds of concepts? *Annual Review of Psychology, 51*, 121–147.
Murphy, G. L. (2002). *The big book of concepts*. Cambridge, MA: MIT Press.
Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General, 115*, 39–57.
Osherson, D. N., Smith, E. E., Wilkie, O., López, A., & Shafir, E. (1990). Category-based induction. *Psychological Review, 97*, 185–200.
Rogers, T. T., & McClelland, J. L. (2004). *Semantic cognition: A parallel distributed processing approach*. Cambridge, MA: Bradford Books.
Rosch, E. (1977). Human categorization. In N. Warren (Ed.), *Studies in cross-cultural psychology* (Vol. 1, pp. 1-49). London: Academic Press.
Rosch, E., & Mervis, C. B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology, 7*, 573–605.
Rosch, E., Mervis, C. B., Gray, W. D., Johnson, D. M., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology, 8*, 382–439.
Ruts, W., De Deyne, S., Ameel, E., Van Paemel, W., Verbeemen, T., & Storms, G. (2004). Flemish norm data for 13 natural concepts and 343 exemplars. *Behavior Research Methods, Instrumentation, and Computers, 36*, 506–515.
Ruts, W., Storms, G., & Hampton, J. (2004). Linear separability in superordinate natural language concepts. *Memory and Cognition, 32*, 83–95.
Shafto, P., Kemp, C., Mansinghka, V., Gordon, M., & Tenenbaum, J. B. (2006). Learning cross-cutting systems of categories. In *Proceedings of the 28th annual conference of the cognitive science society.*
Sloman, S. A., & Malt, B. (2003). Artifacts are not ascribed essences, nor are they treated as belonging to kinds. *Language & Cognitive Processes, 18*, 563–582.
Smith, E. E., & Medin, D. L. (1981). *Categories and concepts*. Cambridge, MA: Harvard University Press.
Smith, J. D., & Minda, J. P. (2000). Thirty categorization results in search of a model. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 26*, 3–27.
Smits, T., Storms, G., Rosseel, Y., & De Boeck, P. (2002). Fruits and vegetables categorized: An application of the generalized context model. *Psychonomic Bulletin and Review, 9*, 836–844.
Sneath, P. H. A., & Sokal, R. R. (1973). *Numerical taxonomy*. San Francisco, CA: Freeman.
Storms, G. (2001). Flemish category norms for exemplars of 39 categories: A replication of the Battig and Montague (1969) category norms. *Psychologica Belgica, 41*, 145–168.
Storms, G. (2004). Exemplar models in the study of natural language concepts. In B. H. Ross (Ed.), *The psychology of learning and motivation*. New York: Academic Press.

Storms, G., & De Boeck, P. (1997). Formal models for intracategorical structure that can be used for data-analysis. In K. Lamberts & D. Shanks (Eds.), *Knowledge concepts and categories* (pp. 439–459). London: UCL Press.
Storms, G., De Boeck, P., & Ruts, W. (2000). Prototype and exemplar based information in natural language categories. *Journal of Memory and Language, 42*, 51–73.
Storms, G., De Boeck, P., & Ruts, W. (2001). Categorization of unknown stimuli in well-known natural concepts: A case study. *Psychonomic Bulletin and Review, 8*, 377–384.
Storms, G., Van Mechelen, I., & De Boeck, P. (1994). Structural analysis of the intension and extension of semantic concepts. *European Journal of Cognitive Psychology, 6*, 43–75.
Sutcliffe, J. P. (1993). Concept, class, and category in the tradition of Aristotle. In I. Van Mechelen, J. A. Hampton, R. S. Michalski, & P. Theuns (Eds.), *Categories and concepts: Theoretical views and inductive data analysis*. London: Academic Press.
Tversky, A. (1977). Features of similarity. *Psychological Review, 84*, 327–352.
Van Mechelen, I., Bock, H.-H., & De Boeck, P. (2004). Two-mode clustering methods: A structural overview. *Statistical Methods in Medical Research, 13*, 363–394.
Van Mechelen, I., De Boeck, P., Theuns, P., & Degreef, E. (1993). Categories and concepts: Theoretical views and data analysis. In I. Van Mechelen, J. A. Hampton, R. S. Michalski, & P. Theuns (Eds.), *Categories and concepts: Theoretical views and inductive data analysis*. London: Academic Press.
Zeigenfuse, M., & Lee, M. (2010). Finding the features that represent stimuli. *Acta Psychologica, 133*(3), 283–295.