Chapter 4 Conceptual Spaces, Features, and Word Meanings: The Case of Dutch Shirts

Joost Zwarts

Abstract This paper explores how a conceptual space for the representation of word meanings can be constructed and visualized for one particular domain, namely Dutch words for different types of shirts. It draws on earlier empirical corpusbased research that has identified different features for uniquely describing each of these types and different ways in which they are lexically described in fashion magazines. The present study defines a metric that makes it possible to construct a feature-based space in which the extension of each of the Dutch shirt terms can be visualized and in which it is possible to study the distribution of words and the validity of different constraints on that distribution: conjunctivity, convexity, connectivity, coherence, and centrality. Although the paper concludes that definite conclusions about these constraints are only possible on the basis of more complete lexical datasets, it demonstrates the potential of the conceptual space approach for studying word meanings.

4.1 Conceptual Spaces and Semantic Maps

One way of doing lexical semantics is by studying particular meaning domains as *conceptual spaces* or *semantic maps* (see Gärdenfors 2000; Haspelmath 2003, respectively, for general overviews). The idea is that a domain consists of a set of values as points, geometrically structured in a particular way, with lexical categories (extensions) as regions. The geometrical structure of the domain could be assumed to be universal, but languages divide it up in different ways. One well-known example is the color space, with its dimensions of hue, saturation, and brightness (Kay et al. 2009). Another example is the graph of functions of indefinite pronouns (Haspelmath 1997), as shown in Fig. 4.1.

F. Zenker, P. Gärdenfors (eds.), Applications of Conceptual Spaces,

Part of this paper was presented at the workshop *Conceptual Spaces at Work*, Lund, May 24–26, 2012.

J. Zwarts (🖂)

Department of Languages, Literature and Communication, Utrecht University, Trans 10, 3512 JK Utrecht, The Netherlands

e-mail: J.Zwarts@uu.nl

[©] Springer International Publishing Switzerland 2015

Synthese Library 359, DOI 10.1007/978-3-319-15021-5_4



Fig. 4.1 Graph of functions of indefinite pronouns

The functions that indefinite pronouns can have are organized in a graph, with more similar functions closer to each other. Indefinite pronouns with *some* (e.g. *something*), *any* (e.g. *anything*), and *no* (e.g. *nothing*) correspond to contiguous sets of functions on the graph.

In one form or another this 'spatial' approach has been used for such diverse domains as modal verbs like *may* and *can* (Van der Auwera and Plungian 1998), container nouns like *jar* and *bottle* (Malt et al. 1999), motion verbs like *climb* and *crawl* (Geuder and Weisgerber 2002), adpositions like *in* and *on* (Levinson and Meira 2003), verbs of cutting and breaking (Majid et al. 2008), and case markers like the dative and accusative (Grimm 2011), to name just a few.

A conceptual space consists of a set of 'meanings' (like colors, referential functions, modalities, pictures of containers, pictures of spatial relations, video clips of cutting and breaking events, bundles of semantic properties) and some mathematical structure defined over that set.¹ This structure can be a discrete graph (like Fig. 4.1) or it can be a continuous metric (like the color space), but the idea is always that meanings that are closer together in the conceptual space are more similar, like the specific or negative functions in Fig. 4.1.

There are at least three ways in which one can construct such a similarity space for a domain:

- The lexical way: Meanings are closer if speakers use the same lexical item for them more often across languages.
- The psychological way: Meanings are closer if human subjects judge them to be more similar, non-linguistically.
- The semantic way: Meanings are closer together if a semantic analysis treats them as more similar.

The lexical way is the one most traveled. We find it in the graph-based semantic maps of typological linguistics (Haspelmath 2003) and in the statistical approaches

¹Note that the term 'meaning' is used here in a very broad sense, encompassing both specific referents and general functions, and both non-theoretical and theoretical notions of meaning.

of Levinson and Meira (2003) and Majid et al. (2008). For example, in the space of Fig. 4.1, the meanings 'conditional' and 'comparative' are close together because languages tend to use the same word for those meanings, like English does with *any*. The psychological way of deriving a conceptual space can be seen in Malt et al. (1999) and in much other work that uses pile sorting or other ways to derive non-linguistic similarity judgments of stimuli. Van der Auwera and Plungian (1998), Geuder and Weisgerber (2002), Grimm (2011), and Gärdenfors et al. (2012) build their spaces on semantic considerations, that is, on an analysis of the functions or referents that are covered by the expressions that they study.

In Zwarts (2010), I argued that conceptual spaces need to be approached from different, complementary angles, especially if we want to study constraints on lexical categorization. If we want to understand how words can be meaningfully used and learned, then we need to understand what can and cannot be a word meaning. Gärdenfors (2000) and Haspelmath (2003) independently argued for *geometric* constraints on meaning. According to Gärdenfors, meanings are *convex*, while in Haspelmath's semantic maps they are *contiguous* (connected, in graph theoretical terms). Such constraints also play an important role in the modal map of Van der Auwera and Plungian and the case map of Grimm. Obviously, if we want to test whether the regions corresponding to words satisfy certain constraints, then those constraints should not themselves be part of the recipe for making the space, as in the lexical approach, but the space should be constructed independently of the lexical items that are distributed over it. Only then can we test the hypothesis that word meanings are convex or contiguous.

In a sense, this paper is an experiment. It starts with a lexical domain described in Geeraerts et al. (1994) (henceforth GGB) for which three things were given: (i) a set of referents (244 types of shirts), (ii) a set of words from one language applying to those referents (seven Dutch words), (iii) an analysis of each of these referents in terms of properties (shape, fabric, fastening, ...). The set of referents is too big in relation to the set of words to follow the lexical route to a conceptual space. There are only seven words and 244 referents, so we cannot induce the similarity relations of the referents on the basis of those words. But suppose we would use the feature analysis of the referents that GGB provide, what kind of space would we get, with what sort of 'shape', and how is the distribution of the words constrained by the geometry of the space? Answering these questions is what I set out to do in this paper. Although the results are not entirely conclusive (because of the nature of the data, as we will see), the approach looks promising because of the way it explicitly links constraints of categorization to semantic features, making them testable, in principle.

This paper is structured as follows. After introducing the domain of shirts in Sect. 4.2, I will define the 'shirt space' in Sect. 4.3. This space has a particular 'shape' that shows two major clusters of shirt types (Sect. 4.4). I then discuss the status of four different categorial constraints in this feature-based shirt space: conjunctivity (Sect. 4.5.1), convexity (Sect. 4.5.2), connectivity (Sect. 4.5.3), and coherence (Sect. 4.5.4).

4.2 Dutch Shirts

GGB report the results of research done at the University of Leuven, Belgium, between 1990 and 1993, into the nature and origins of lexical variation. A total of 9,000 occurrences of clothing terms were collected from magazines in Dutch, published in Belgium and the Netherlands. In each case the occurrence of the term was accompanied by a picture that showed an instance of the item. These pictures were used to make componential analyses of the referents of clothing terms.

For example, one of the referents encountered has feature decomposition [52131] and is found labeled by words like *spijkerbroek* and *jeans*, both of which mean 'jeans'. Each of the five digits of [52131] represents a value on a particular dimension:

Length: 5 (down to the ankles) Width and cut: 2 (straight cut, neither tight nor wide) End of legs: 1 (no special features) Material: 3 (denim) Details: 1 (strengthened by metal buttons)

There might be other referents, with different features that are also labeled *spijkerbroek* or *jeans*. This means that every clothing term has an extension that consists of referents, each uniquely described by a set of discrete features. It is important to realize that [52131] is not the traditional semantic decomposition of a *word meaning*, but the analysis of one particular *referent*.² GGB were interested in determining the prototype structure or family resemblance structure of clothing categories, not necessarily their classical definition in terms of necessary and sufficient features, which might not always exist.

There are not many details about how the researchers went about to make their feature decomposition and the original pictures are not included in the book. There are probably different ways in which one can analyze a set of clothing items into features, and certain decisions are made that would need further motivation, but nevertheless, I am assuming for now that their analysis into features is at least a plausible way to represent what the referents are like, coming back to possible shortcomings along the way and in the conclusion.

I focus on one particular subdomain from their book, which involves 'shirts, tshirts, blouses; garments covering the upper part of the body, made of light material, constituting the first layer of clothing above the underwear' (p. 22). All the data are taken from their book (p. 129–133). They present a total of 244 configurations (referent types), that were analyzed for features of shape, length, fastening, fabric,

 $^{^{2}}$ A referent like [52131] is a *type* of object that corresponds to many different *tokens* that share these five properties, but differ in a lot of other properties.

collar, neckline, position of the buttons, and gender.³ There were 7 Dutch terms found for these items, which are all covered more or less by the English noun *shirt*, each of which has as its extension a proper subset of the 244 shirt types. The following glossing is only a very rough approximation:

blouse 'blouse', hemd 'shirt', overhemd 'dress shirt', overhemdblouse 'shirtwaist', shirt 'shirt', t-shirt 't-shirt', topje 'tank top'

On the basis of the data they collected for these terms, GGB demonstrated that this field has a 'non-classical, un-mosaic-like character' (p. 134), without 'sharp divisions between the individual items within the field' (p. 118). The terms are overlapping to a large extent and they are not hierarchically ordered: *overhemd*, for instance, is not a hyponym of *hemd*, and *t-shirt* not of *shirt*.

The features with their values are as follows, each given with an example:

Shape: 1,2,3,4 (e.g. 4 = covering trunk and arms) Length: 1,2,3,4 (e.g. 2 = tucked into skirt or trousers) Fastening: 1,2,3,4 (e.g. 3 = full fastening) Fabric: 1,2,3,4 (e.g. 1 = smooth, cottonlike) Collar: 1,2,3,4,5 (e.g. 3 = soft collar) Neckline: 1,2,3,4,5,6 (e.g. 3 = round neckline) Position of the buttons: 1,2,3 (e.g. 2 = left) Sex: v,m (e.g. v = female)

GGB have used discrete feature values for dimensions that are truly discrete, like the position of buttons (left, right) or the gender of the shirt (male, female), but also for dimensions that are really continuous, like length, by partitioning this continuous scale into a small number of intervals.

Each particular shirt corresponds to a string of feature values. Referent number 3, for instance, corresponds with the string [4231332v]. Following a common linguistic practice, I put square brackets around a feature bundle. I refer to referents with S1, S2, etcetera, in order not to confuse them with feature values. The ordering of the values of a feature is not significant, although in some cases the choice of integers is not entirely arbitrary. For instance, the other values of Shape are 1 =covering trunk below shoulders, 2 = covering trunk and shoulders, leaving arms uncovered, and 3 = covering trunk, shoulders and upper arms. I will treat the values of all features as unordered.

Because we are working with strings, it is possible to pick out classes of shirts with regular expressions in the usual way, as a useful notion. The full stop (.) is used as the wildcard for any value of a feature, the vertical bar (|) for alternative values, and the dash (-) for a range of values. The strings of features can now be used to define our 'shirt space', which is the topic of the next section.

³The book actually gives 246 configurations, but there are two pairs with the same feature profile, so I counted both of these pairs as one referent. Therefore the configurations numbered 67 and 68 are missing in my list.

4.3 Shirt Space

The idea of a conceptual space is to represent semantic similarity of the elements of a domain in terms of spatial distance. For our situation this means that we have to determine distances between shirts on the basis of their feature makeup. For this I use a very simple metric, namely the *Hamming distance*.⁴ If we consider only strings of equal length, then the Hamming distance between two strings is simply the number of positions at which they differ. For instance, the distance between configuration S1, characterized by the string of features [3231432v], and configuration S2, with the string of features [3431412v] is 2, because they differ only in the second and sixth position. For short, I will use d(x,y) for the distance between two feature bundles x and y in a set of such bundles S. Given the way distance is defined, we have the following properties for all configurations x, y, and z in S:

 $d (x, y) \ge 0$ d (x, y) = 0 iff x = yd (x, y) = d (y, x) $d (x, y) + d (y, z) \ge d (x, z)$

This makes our set of shirts a metric space, in which similar shirts are closer together (see Gärdenfors 2000 for the metric properties of conceptual spaces).

Note that this similarity metric is very simple. It does not take into account that some features (like Sex) have two values, while other features (like Neckline) have six. The distance between a female and male version of a shirt is just as great as the distance between a shirt with a round or rectangular neckline. In this similarity metric all the features have the same weight, to keep things simple, but also because it is not straightforward to determine what the weights would have to be for this set of data.

In a sense, this approach follows the opposite direction from approaches that start with lexical or similarity judgment data and perform multidimensional scaling or a similar statistical operation to derive the features. By starting with features, instead, the dimensions of the space are already given, because each feature with its range of values can be seen as representing a dimension. We usually think of a dimension as a continuous scale, but here the dimensions only have a few discrete and unordered values. It is on the basis of these 'a priori' given features/dimensions that a metric space is defined over the set of referents. Multidimensional scaling works in the opposite direction. It starts with a high-dimensional metric space for a set of referents (based on how often referents are named by the same word across languages, for instance), and then extracts a few dimensions that best represent this high-dimensional space.

⁴For more sophisticated similarity measures based on features see Tversky (1977).

This distance metric d can be used to define a notion of betweenness b.

For all distinct x, y, z in S, y is between x and z if and only if d(x,y) + d(y,z) = d(x,z).

So, for instance, referent S107 [3431432v] is between S1 [3231432v] and S2 [3431412v], because d(S1,S107) = 1 and d(S107,S2) = 1 and d(S1,S2) = 2. We can say that two distinct elements are *adjacent* when there is no element between them:

For all distinct x and z in S, x is adjacent to z if and only if there is no y in S such that y is between x and z.

Referent S1 and S2 are not adjacent to each other, but they are both adjacent to S107. Adjacent referents can have a distance greater than 1: S1 and S102 [3231431m] have a distance of 2, but they are adjacent because there are no referents with values [3231432m] or [3231431v].

Finally, we can define the notion of a *path* in a space *S*:

For all distinct x and y in S, a path from x to y is a sequence of distinct elements x,y_1,\ldots,y_n,z in S with $n \ge 0$, such that every two subsequent elements are adjacent.

For example, one path from S1 [3231432v] to S4 [4231411v] is the sequence S1,S6,S109,S4, where S6 is [4231432v] and S109 is [4231412v]. In this case, each step of the path corresponds with one feature difference. There can of course be more paths between two referents and the distances between the referents on the path can be two or more.

With this feature decomposition, we can now try to address the questions that we posed earlier. What kind of constraints can we find on lexical regions defined over a feature-based space? But first we look at the shape of the space itself that is defined by the features.

4.4 The Shape of the Space

With the number of features and values presented in the preceding section, a total of 38,400 possible items can be defined. However, there are only 245 actual distinct items in this data set, which means that we are dealing with an irregularly shaped semantic space, in which some areas are more populated than other areas. This might have partially to do with the inevitable restrictions of sampling: certain types of shirts might exist but were simply not found in the magazines used for the sample. Also, the sample was biased to women's clothing, because of the nature of the magazines. But there are definitely also some real constraints on the space.

• Cultural constraints, having to do with the kind of shirts that are worn by men and women. For example, those shirts that only cover the trunk below the shoulders are exclusively worn by women (in other words, there are no referents [1.....m] in the database, where 1 means 'only covering the trunk below the shoulders' and m is 'worn by men').

- Constraints of a more logical nature. If a shirt is specified as having no fastening, then the buttons are neither right nor left (so there are no referents with the feature description [..1...(1|2).]).
- Physical (or 'technical') constraints. A wide neckline, for instance, does typically not allow for a collar. (Of the twelve [.....6..] shirts, eleven are [....16..] and one is [....36..]). 6 is the wide neckline feature, 1 is the no collar feature.

In this way, one might expect the shirt space to have an irregular shape, in some general sense like what has been argued for the color space (Regier et al. 2009), which has different saturations for different hue-lightness combinations. However, unlike the color space, the shirt space does not owe its irregularity to perceptual factors, but to other factors. Unfortunately, there is not enough information in GGB to pursue this topic more. Therefore, we have to look at more general ways to look at the shape of the space.

There are different ways in which one can spatialize and visualize a space like this. The graph in Fig. 4.2 shows all the shirts, but draws only edges between two shirts if they differ in exactly one feature.⁵ Most of the shirts are connected in this way, but some are floating around unconnected, simply because their distance to other shirts is greater than 1. They are randomly placed by the drawing program.

Notice that we do not get an explicit representation of the dimensions of the space in this way. These dimensions are spread out, in a sense, across the graph, as we can see when we give the vertices of the graph a color that represents with a particular feature value. For instance, in Fig. 4.3 the four different values of the feature Shape are distinguished by color in the following way (descriptions taken from GGB, p. 129):

Blue: Covers the trunk below the shoulders ([1.....]).

Red: Covers trunk and shoulders but leaves the arms uncovered ([2......]).

Green: Covers trunk, shoulders and upper arms, but leaves the lower arms uncovered ([3.....]).

Yellow: Covers the trunk and the arms ([4.....]).

We can only make the feature dimensions more explicit if we take 'cross sections' of the shirt space. A two-dimensional cross section might consist of the features Sex (with values m, v) and Shape (with values 1, 2, 3, 4), as shown in Fig. 4.4. The values of Shape are ordered from covering less to more of the upper body, which happens to correspond to the ordering of the corresponding integers. The combination [1....m], here abbreviated as m1, as we already saw (there are no shirts for men that only cover only the trunk below the shoulders).

Two clusters present themselves in the space of shirts in Figs. 4.2 and 4.3. When we inspect how the features of those two clusters differ, then we actually see that those clusters correspond roughly to two types of shirts. In Fig. 4.5 two broader types of shirts are indicated. Green corresponds to the more formal type

⁵The graph is drawn by the graphviz software package, using the neato command, which creates a layout that approximates distances in the graph.



Fig. 4.2 Shirts with one feature difference linked

of dress shirts, with a full button fastening and a collar ([...3.[2-5]...]), while red corresponds to the less formal type of shirt, lacking buttons and collar ([....1.3.]).⁶

 $^{^{6}}$ Notice that graphviz does not draw the same graph with a fixed orientation, which is why the graph in Fig. 4.5 is rotated.



Fig. 4.3 Distribution of the shape feature

The two clusters seem to have a kind of feature-based identities. To what extent is that reflected in the naming patterns? Unfortunately, we can not represent all the seven shirt nouns together in one graph, because they overlap each other and hence do not partition the space. Figure 4.6 shows how three of the nouns distribute their extension over the space. Yellow is *topje*, red is *t-shirt*, blue is *overhemd*, orange is used for the shirts that are labeled both *topje* and *t-shirt* and purple is used for the shirt that is both named *t-shirt* and *overhemd*.



Fig. 4.5 Two types of shirts and their correspondence to clusters

As we can see, the names correspond quite well to the clusters. The items called *overhemd* (coloured blue) cluster in the more formal area of the space, while the items called *t-shirt* (red) and *topje* (yellow) cluster in the less formal area. Some members of *topje* are found in the *overhemd* cluster, which on closer inspection turn out to be ones that have a full fastening with buttons.



Fig. 4.6 Overhemd, topje, and t-shirt in the space of shirts

After this exploration of the structure of the space, we now turn to different ways in which the shirt categories might be constrained in terms of the underlying, feature-based space.

4.5 Constraints on the Categories

4.5.1 Conjunctivity

One theoretical possibility is that the shirt categories are classically definable by means of necessary and sufficient conditions, that is, as a conjunction of one or more features. We have already seen such categories in the previous section. For example, when we are talking about shirts with a collar and a full fastening with buttons or shirts that do not have a fastening, buttons or a collar, then we are using conjunctive definitions.

With a collar and a full fastening with buttons: [..3.[2-5].[1-2].] Without a fastening, collar or buttons: [..1.1.3.]

For every feature, there is either no specification (.) or a range of one or more values (like [2-5]). A disjunctive definition would be a class of shirts that are either collarless *or* that have buttons, i.e. $[....1..] \vee [.....[1-2].]$, with two regular expressions.

As we know, conjunctive definitions make a lot of intuitive sense and they play an important role in certain domains (see Hage 1997). At the same time, one of the reasons that GGB undertook their empirical study of clothing terms is that they wanted to demonstrate that not all clothing terms allow for a conjunctive definition, but that prototype and family resemblance structure play an important role. In fact, it is not possible for the set of data that we have from GGB to come up with conjunctive characterizations of any of the shirt nouns. In order to cast the net wide enough to cover all the positive referents of a noun, we get also referents in our net that are not attested with that noun. For example, when we look at items S6 ([4231432v]) and S133 ([4231432m]) in the extension of *overhemd* we see that the sex for the person for whom the shirt is meant should not matter. So for every other female referent of *overhemd*, the corresponding male version should also be in the extension (if it exists). In other words, *overhemd* should be unspecified for sex. However, we find pairs of referents differing only in this feature and which are not both in the extension of *overhemd*:

S136: [4234411m] (*overhemd*) versus S224: [4234411v] (*blouse*) S111: [4232431m] (*hemd*) versus S112: [4232431v] (*overhemd*)

However, I hesitate to draw definite conclusions from this about whether shirt categories might be conjunctively defined. First of all, we have to realize that, however rich the data are, they might still not be exhaustive. It could very well be that with a bigger corpus of magazines, *overhemd* would have been attested for referent S224 and for referent S111. Asking mother tongue speakers of Dutch how acceptable they find *overhemd* for the relevant pictures might also have given different results. This point is also relevant for other constraints that we discuss in the next three sections. Second, it could be that *overhemd* has a classical, conjunctive definition, but that part of its extension is blocked by other words (*blouse* or *hemd*) that are more appropriate for that part. This is also something that can only be ascertained by a very large corpus or by deliberate elicitation.

Nevertheless, as far as the data go, the conclusion must be that conjunctivity is not a constraint that holds of word meanings in this domain.

4.5.2 Convexity

The constraint of convexity has been proposed by Gärdenfors (2000) as a constraint on natural properties, conceived as regions in a particular integrated domain. Even though it is not clear whether shirt names should be seen as referring to natural properties, still I believe it is worthwhile to investigate whether the extensions of the nouns are convex given the underlying feature space. The definition of convexity is as follows:

A subset *C* of a space *S* is convex if and only if for every *x* and *y* in *C*, all points between *x* and *y* are also in *C*.

Intuitively, convexity would make a lot of sense for shirt categories. If we take the shirts number S30 = [2331132v] and S35 = [3311133v], which are both called *t-shirt*, then what they have in common can be written as [.3.113.v]. The idea of convexity is that every referent that has these specific feature values and shares its other features with either S30 or S35, should also be called *t-shirt*. The referents that we find between S30 and S35 are S44 = [2331133v], S81 = [2311133v]), and S188 = [3331132v]. Of these three, only S44 is called *t-shirt*, while S81 is called *shirt*, and S188 *blouse*. So, already in this randomly picked example, convexity does not seem to work, but we have to look at it in a more general way.

We can get an idea of the extent to which categories are convex in this domain by using the notion of a *convex hull*, the closure of a set under betweenness.

In a space S, the convex hull H of set E is the union of E with those elements of S that are between members of E.

Here is a small example. Suppose that $E = \{S1 \ [3231432v], S2 \ [3431412v]\}$. There are two referents between S1 and S2, namely S107 [3431432v] and S181 [3231412v]. As a result, $H = \{S1, S107, S181, S2\}$. In a sense, we make the convex hull by filling up the 'hole' between S1 and S2.

It turns out that the convex hull is quite a bit bigger than the extension, for all of the nouns. In other words, there are quite a lot of 'holes' in the extensions. The numbers are as shown in Table 4.1.

Table 4.1 Non-convex	Noun	Extension	Convex hull				
complements of categories	blouse	116	210				
	hemd	30	127				
	overhemd	56	132				
	overhemblouse	7	37				
	shirt	25	141				
	t-shirt	37	120				
	topje	35	100				

Every item that is in the convex hull of a category, but not in its extension is a counterexample against convexity. This clearly shows that these categories are not convex in the sense defined here. Again, different interpretations of this result are possible. As we already saw in the previous section, the corpora might not have provided enough naming data, thereby creating these 'holes'. Another response could be that it is not the category as a whole that is convex, but that it is only convex in certain dimensions. For example, clothing categories might be convex on the dimensions of shape. This seems also more in line with the position in Gärdenfors (2000). Yet another response might be that convexity is too strong a constraint here. What we need instead is *star convexity*. The category is then organized around a central point (a prototype) such that every referent in the category can be connected with this prototype by a line that is entirely within the category. This idea would be in line with a historical development of categories, from a prototype in different directions of similarity, maybe like the *chaining* in Malt et al. (1999). I leave it to further research to investigate this possibility.

4.5.3 Connectivity

While convexity is the constraint in Gärdenfors (2000), what we find in the semantic maps of Haspelmath (2003) is a weaker property of contiguity or, rather, connectivity in graph theoretical terms. The idea is that a word forms a connected subgraph of an underlying conceptual space. Maybe the shirt nouns are connected in this sense on the feature-based graph?

In order to investigate this, we need to construct a graph with the right kind of connections, for which adjacency seems appropriate. Remember that two meanings are adjacent if and only if there are no other meanings between them, which comes close to the kind of relation that underlies connectivity in semantic maps. With this relation, most of the shirt categories seem completely connected, like *overhemd* in Fig. 4.7, for instance, but unlike *hemd* in Fig. 4.8, which has two members (colored red) which are not adjacent to any other member. These diagrams show only the members of the categories together with their adjacency structure.





Fig. 4.7 The connectivity of overhemd

The problem is however that the kind of connectivity that adjacency gives us is too liberal. The reason is that every referent has quite a lot of adjacent neighbours. The number ranges from 5 to 63, with an average of a little bit over 17. This means that even if we would form an arbitrary set of elements, then the chance of each member to be adjacent to at least one other member of that set is quite big. The reason that our shirt space behaves this way might be that the theoretical space of possible types is quite large, as we saw, but because there are many gaps in the space, there are adjacencies over long distances, making this a quite tightly knit space, as a whole. An alternative way of defining adjacency would be in terms of Hamming neighbors, referents that differ only in exactly one feature. However, because of the sparsely filled space, this version will tend to underestimate connectivity.

Let me turn to the last possibility.



Fig. 4.8 The connectivity of hemd

4.5.4 Coherence

A common idea in the literature is that categories are 'coherent' or 'compact', meaning that they maximize the within-category similarity and minimize the across-category similarity (Rosch and Mervis 1975; Tversky 1977; Regier et al. 2009, among others). These last authors even showed that the partitions of the colour space by languages are near-optimal in the sense that for a whole partition the well-formedness is higher than for alternatives that are derived by rotating the color space with respect to a particular set of terms. Such rotations yield alternative terminological systems that are closely related to the original. A full exploration of this idea for the shirt space would go about defining rotations of the shirt space to derive close related alternative categorizations.

Nevertheless, 'rotation' in a looser sense is an easy way to shift the extension of a category to a more random alternative category. Suppose we shift a given category to a new category by shifting all its referents up in the list by a particular number,

Noun	No shift	+10	+20	+30	+40	+50
blouse	4	4.2	4.3	4.4	4.5	4.5
hemd	3.6	4.1	3.9	3.9	4	4.6
overhemd	3.4	4.5	4.4	4.2	4.1	4
overhemdblouse	2.5	3	2.8	3.4	3.4	2.8
shirt	4.7	4.8	4.2	4	3.9	3.5
t-shirt	3.5	3.8	4.3	4.7	4.7	4.3
topje	3.7	4.1	3.8	3.7	3.9	4.4

Table 4.2 Average distance within categories and their shifted versions

going by the numbering of GGB.⁷ We can compute the average distance among the members of a set of shirts in the following way:

If *S* is a set of *n* items with a distance metric *d*, then the average distance D(S) of *S* is given by $\sum d(x,y)$ for each pair $x,y \in S$, divided by $\frac{1}{2n(n-1)}$.

What we can see in Table 4.2 is that the average distance of categories that are shifted (over five different distances) is generally higher than the average distance of the original category, which supports the idea that categories have some sort of coherence. However, there is one interesting counterexample, namely *shirt*. The average distance within this category is a bit higher than within the other categories and some of the shifts of *shirt* make the category actually more coherent it seems. One interpretation might be that *shirt* does not have a well-established shared meaning across different users and therefore lacks in coherence.

Let me finish this section by showing how we can get a visual impression of coherence when we display the categories in a graph that also respects the distances between the nodes. In this display, multidimensional scaling is used as a model for approximating the distances between the nodes. As we can see in Fig. 4.9, the region corresponding to *hemd* is fairly coherent, with only one clear 'outlier'. The red nodes are the disconnected ones, in the sense described earlier. Figure 4.10 shows that *shirt* distributes over the space in a much less coherent way.

It is also possible to identify members of a category that minimize the average distance to other members, and are central in that sense. For the category *hemd*, referent S4 has the smallest average distance to other members. One might want to say that S4 is like the prototype of this category. It should have features then that are more typical of this category and this actually turns out to be true. S4 has the profile [4231411v] and when we look at the feature values that occur most frequently with referents of *hemd*, as shown in Table 4.3, then we can see that most frequent values for each feature (highlighted by boldface) are exactly the feature values of the referent that is spatially central, thereby reflecting the redundancy structure of the category as a whole (Rosch and Mervis 1975).

⁷There is no system in the list, apart from the fact that sometimes a contiguous range of items in the list seem to belong to the same category.



Fig. 4.9 The coherence of hemd

4.6 Conclusion

This paper has explored how we can construct a conceptual space on the basis of a given set of features and thereby study different properties of linguistic categories. For this particular domain of shirts and Dutch data set, the patterns might be linked to what has been found in another domain of artefacts, namely (household) containers (e.g. Malt et al. 1999, 2010), where the link between the perceived features and the linguistic labelings seems fairly 'loose' and strongly influenced by language-specific and culture-specific factors.

We have also seen that the extensions of shirt nouns show coherence in terms of the underlying feature space, but that they do not show the kind of conjunctivity or convexity that we would expect if they were based more directly on concepts (like 'long-sleeved shirt with a stiff collar with full fastening with buttons'). It is conceivable that these categories are not only held together on the basis of the underlying space, but also by conventions that cause referents to belong to the same category, even if they are not spatially related in one way or another. We know that metaphorical and metonymical mappings can extend the application of a term in a 'non-local' way and similar mechanisms might be at work in this domain



Fig. 4.10 The lack of coherence of shirt



Feature	1/v	2/m	3	4	5	6
Shape	1	1	7	21	-	-
Length	0	14	11	5	-	-
Fastening	1	0	29	0	-	-
Fabric	20	4	1	5	-	-
Collar	2	0	3	25	0	-
Neckline	16	3	11	0	0	0
Buttons	16	13	1	_	_	-
Sex	19	11	-	-	-	-

too. However, there are several considerations that need to be kept in mind when evaluating the categorial constraints, having to do with the nature of the data and with the nature of the constraints.

The data of this study come from a corpus of fashion magazines. Although GGB made sure that the corpus was saturated in the sense that a bigger corpus would not have contained more *word types* or *referent types*, still a bigger corpus might have

given more *naming relations*, that is, applications of words to referents. If a referent r is not named by noun n in the corpus, then this does not allow us to conclude that certain strong constraints, like convexity, do not hold. Elicitation of naming relations directly from native speakers seems a better way of collecting the relevant data, but even here care should be taken that acceptability of a noun n for a referent r is tested exhaustively. Even in the experimental studies of naming there is a tendency to go for the most frequent label of a referent, which is not a good measure if we want to study strong constraints on lexical categories.

Another problem of corpus data for studying general constraints on categorization is that they might be constituted of rather different language varieties, at the level of idiolects, sociolects, or dialects. The corpus of GGB was deliberately composed in such a way that such variation could be studied. However, the result could be that a noun lacks a particular property (like convexity) because its extension in the data set is the union of two or more different uses of that noun, each with their own extension. While coherence might still hold of such an aggregated extension, properties like conjunctivity or convexity are better treated as constraint on the categorizations of individual language users.

As we saw, the elements of the domain are all very close together, leading to an overall high level of connectedness. It seems more likely that features or feature values do not all contribute to the structure of the space and the constitution of categories in the same way. Certain features are more salient than others and will play a greater role in categorization. As Tversky (1977) already showed, features can be weighed, and this might affect the results and conclusions in important ways. The question is then, of course, how one could define weights for the features that define a shirt.

This question is part of the more general question of how the features can best be defined and motivated. Which features are used by humans to categorize shirts and what determines the salience of those features in their perception and categorization? What kind of values can a feature take, ranging from binary (male/female user) to multidimensional and continuous (the shape)? How do those feature values affect the shape of the conceptual space and the distances between referents in that space? I have shown that a conceptual space and spatial constraints can be meaningfully defined on the basis of discrete features, but the results might be different if continuous dimensions are used where possible.

Finally, there are more sophisticated methods and techniques that could be used to study feature-based spaces, such as network analysis and machine learning, and all the statistical and graphical methods that are part of that. In that way, inductive and deductive, discrete and quantitative, theoretical and empirical approaches to word meaning can be more tightly integrated within the general perspective of conceptual space semantics.

Like I said at the beginning, this paper presents an experiment in the exploration of a conceptual space. Although it does not yield definite results about how concepts in conceptual space might be constrained in terms of the underlying features, it does show the usefulness of such an approach, especially if this linguistic approach would be wedded with both computational and psychological methods. Acknowledgements I thank the audience for helpful questions and remarks as well as two anonymous reviewers and Peter Gärdenfors for their comments. The Netherlands Organization for Scientific Research (NWO, grant 360-70-340) and the Swedish Collegium for Advanced Study are also gratefully acknowledged for their support.

References

- Gärdenfors, P. (2000). *Conceptual space: The geometry of thought*. Cambridge, MA: The MIT Press.
- Gärdenfors, P., Warglien, M., & Westera, M. (2012). Event structure, conceptual spaces, and the semantics of verbs. *Theoretical Linguistics*, 38(3–4), 159–193.
- Geeraerts, D., Grondelaers, S., & Bakema, P. (1994). *The structure of lexical variation: Meaning, naming, and context*. Berlin/New York: Mouton de Gruyter.
- Geuder, W., & Weisgerber, M. (2002). Verbs in conceptual space. In G. Kath, S. Reinhard, & P. Reuter (Eds.), Sinn & Bedeutung VI proceedings of the sixth annual meeting of the Gesellschaft für Semantik (pp. 69–83). Osnabrück, Germany: University of Osnabrück.
- Grimm, S. (2011). Semantics of case. Morphology, 21, 515-544.
- Hage, P. (1997). Unthinkable categories and the fundamental laws of kinship. American Ethnologist, 24(3), 652–667.
- Haspelmath, M. (1997). *Indefinite pronouns* (Oxford studies in typology and linguistic theory). Oxford: Oxford University Press.
- Haspelmath, M. (2003). The geometry of grammatical meaning: Semantic maps and crosslinguistic comparison. In M. Tomasello (Ed.), *The new psychology of language* (Vol. 2, pp. 211–242). Mahwah: Lawrence Erlbaum.
- Kay, P., Berlin, B., Maffi, L., Merrifield, W. R., & Cook, R. (2009). *The world color survey*. Stanford: Center for the Study of Language and Information.
- Levinson, S. C., & Meira, S. (2003). 'Natural concepts' in the spatial topological domain Adpositional meanings in crosslinguistic perspective: An exercise in semantic typology. *Language*, 79(3), 485–516.
- Majid, A., Boster, J. S., & Bowerman, M. (2008). The cross-linguistic categorization of everyday events: A study of cutting and breaking. *Cognition*, 109(2), 235–250.
- Malt, B. C., Sloman, S. A., Gennari, S., Meiyi Shi, & Yuan Wang. (1999). Knowing versus naming: Similarity and the linguistic categorization of artifacts. *Journal of Memory and Language*, 40, 230–262.
- Malt, B. C., Gennari, S., & Imai, M. (2010). Lexicalization patterns and the world-to-words mapping. In B. C. Malt & P. Wolff (Eds.), Words and the mind: How words encode human experience (pp. 29–57). Oxford: Oxford University Press.
- Regier, T., Kay, P., & Khetarpal, N. (2009). Color naming and the shape of color space. *Language*, 85, 884–892.
- Rosch, E., & Mervis, C. B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology*, 7, 573–605.
- Tversky, A. (1977). Features of similarity. Psychological Review, 84(4), 327-352.
- Van der Auwera, J., & Plungian, V. (1998). Modality's semantic map. *Linguistic Typology*, 2, 79– 124.
- Zwarts, J. (2010). Semantic map geometry: Two approaches. Linguistic Discovery, 8(1), 377–395.