In: Graham Katz, Sabine Reinhard, and Philip Reuter, eds. (2002), *Sinn & Bedeutung VI*
Proceedings of the Sixth Annual Meeting of the Gesellschaft für Semantik, University of Osnabrück

VERBS IN CONCEPTUAL SPACE*

Wilhelm Geuder & Matthias Weisgerber,
Universität Konstanz

Abstract

In his book 'Conceptual Spaces' Gärdenfors (2000) sets out a framework for the representation of concepts and word meanings which defines semantic features as values on ordered property scales and arranges these scales as independent quality dimensions in a multidimensional coordinate system. The resulting representation of a concept is therefore a region in "conceptual space". In this paper, we will briefly introduce this model, and then explore possible ways of adapting it to verb meanings.

1 Conceptual Spaces: Features, Concepts, and Categories

In this section we shall introduce Gärdenfors' (2000) framework of Conceptual Spaces and illustrate it with some basic examples. The two fundamental notions on which this model is built are features and concepts. *Features* are represented as values on one-dimensional *scales*, where a scale is defined by the following three features: it contains different elements of the same type, it contains each element only once, and there is an ordering relation between the elements. The most straightforward scales are those which can be directly derived from sensory modalities, like sound level or pitch, temperature or the height, width, breadth or weight of an object. To take an example, the feature 'having the temperature w' is represented by a single point on a scale of temperature values w_i which are sorted in increasing order. A *concept* can be defined as a correlation between different features, and each concept allows a certain range of values on each of the scales. When the feature scales that define a concept are used as the axes of a coordinate system, a geometric representation of the concept emerges: it is the (multidimensional) spatial region that is projected from the admissible property values of each axis.

1.1 An Example: the Geometric Representation of the Concept COLOUR

In this section, we illustrate Gärdenfors' model from its application to the concept of colour (this section is entirely based on chapter 1.5 of Gärdenfors 2000). COLOUR is a relatively simple concept, although there is a twist in its representation, as will be seen shortly. Most models describing humans' cognitive representation of colors use three dimensions: saturation, hue and brightness. The descriptive model used by Gärdenfors is the Swedish **N**atural **C**olor **S**ystem by Hård and Sivik (1981).

If we try to define a scale with two extrema—corresponding to the highest and the lowest wavelength that can be perceived—it turns out that these two colors at the extrema are perceived as very similar. This fact about human perception should be reflected as a neighbourhood in the

*original title of the talk: Konzeptuelle Räume: Überlegungen zur geometrischen Repräsentation von Verbbedeutungen model. We therefore need a *color circle* , i. e. a circular ordering of hue values. This is certainly a non-standard case for a property dimension but is required by the fact that the representation needs to be psychologically realistic (see figure 1).

Another property besides HUE which defines colour is SATURATION. In non-technical terms this means for any color *c* we have a scale that ranges from the value ‘fully dominated by color *c*’ to the value ‘fully dominated by background color’. The result can be seen in figure 1.

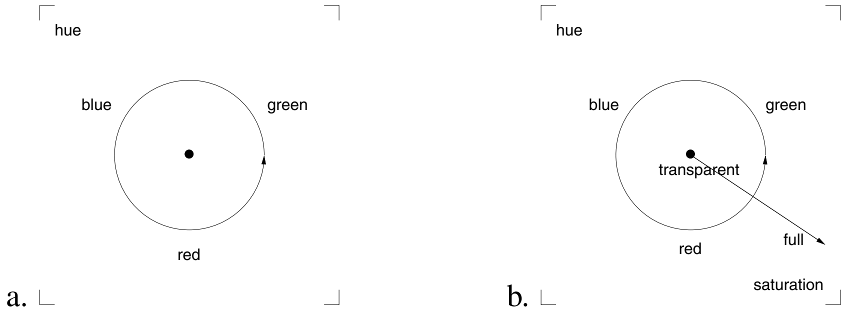

Figure 1: The dimensions of a Color circle, a. ‘hue‘, b. ‘hue’ and ‘saturation’

The third dimension is BRIGHTNESS, which can be defined as a scale between bright white and dark black. Figure 2 shows the color circle with different values for brightness. We chose smaller circles to represent the color circles at the extreme positions of the brightness scale. This reflects the fact that towards either of the extreme values of brightness, the possible ranges of hue and saturation are more and more restricted, and if brightness reaches the value zero, no colour and hue values can be distinguished any more. Mathematically speaking, this points to the fact that the dimensions are not completely independent.

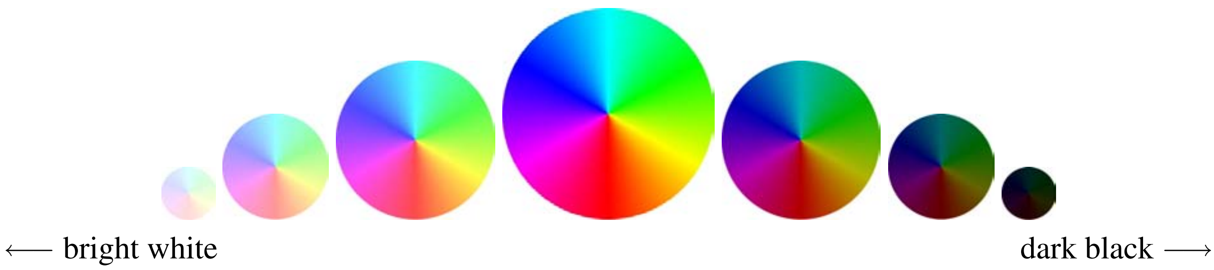

Figure 2: Color circle with different values for brightness, with constant saturation

Now we can combine this information into a three-dimensional coordinate system (see figure 3): In this special case, the graph results from piling up, as it were, an infinite number of color circles along the brightness scale, each one with saturation increasing from the centre outwards. Because of the fact that there are more colour distinctions with medium brightness values than with the extreme ones, the resulting conceptual space for COLOURS is a double cone.

Let us now consider how the model represents, for instance, the conceptual category RED. We find that it is now possible to incorporate the insights from prototype theory into the model: When cutting out a slice in the ‘red’ area of the colour circle, we find a ‘more or less red’ part with a prototypical ‘red’ in its centre and more or less similar ‘reds’ in the periphery. If we perform the same operation with non-focal colours, e. g. ‘dark rose’ or ‘light brown’, we encounter less symmetrical distributions around the prototype. Gärdenfors (2000, p. 73ff.)

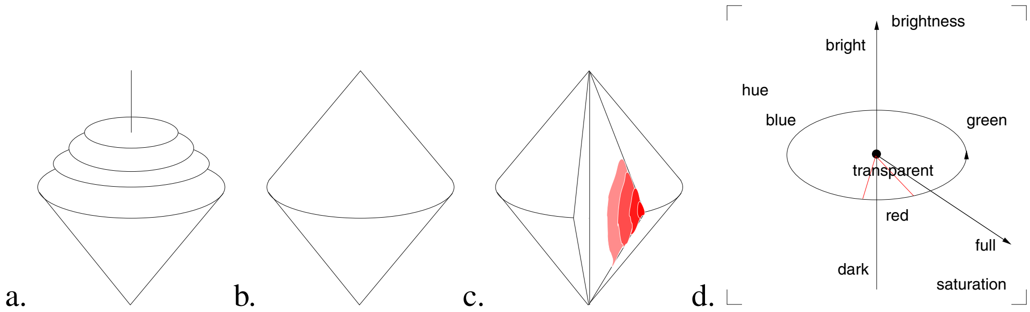

Figure 3: Building the conceptual space for COLOURS and cutting out a slice for 'red'

reports empirical data from Sivik and Taft (1994), who link the cognitive representation of several colors to areas on slices from the double cone.

To conclude, the Natural Color System is a model which represents colours using three feature dimensions, one of them non-euclidic.

1.2 General Properties of the Model

Let us now reflect on the general properties of the geometric model that have become apparent from the above example, and which will bear on its further application to other types of concepts.

First, the whole model is designed to express *similarity relations*, since the values on each property axis are ordered by similarity, and so the distances between two points in the resulting multidimensional space serve as a measure of the overall similarity of the corresponding objects. Second, concepts are not directly decomposed into features; rather, features are always ordered in dimensions.[1] Third, within each dimension, the feature values must be ordered on *scales*. This leads to the necessity of finding a suitable ordering relation. Finding such ordered scales will turn out to be a hard and not always decidable question. In addition, we observe that binary features obtain a special role: finding an ordering relation means deciding which of both elements shall be the low end of the scale; putting both elements on a scale means dividing the scale in two distinct but successive intervals. Although this is well defined, also in a mathematical sense, it may lead to interpretative difficulties. Forth, the model is tailored to representing whole *semantic fields*: There is always a conceptual space projected by the whole range of possible feature values, and this space is then divided up into a certain number of conceptual categories. It does not make very much sense to project a whole such space for the representation of one single item.

In other words, we can say that the model is inherently *contrastive*—in distinction to the symbolistic approach, in which features, and hence semantic categories, can exist "by themselves").

2 Application to Complex Concepts

We now want to return to the question of how conceptual spaces can be used for the representation of more abstract entities, e. g. verb meanings.

[1]An alternative would be to deal with non-ordered *sets* of features, as is done in set-theoretic models like the *feature model* of Tversky (1977).

As typical features of verb meanings, we can note that verbs express relations between participants, and that they encode dynamic situations. We may ask, therefore, how the geometric model can accommodate such properties. To show that this is indeed feasible, we shall first discuss some simple examples of 'relational concepts'. Next, we shall present a sketch of how to encode dynamic patterns, which will be based on the insight that these can be reduced to the case of complex spatial patterns.

2.1 Relational Concepts and Complex Spatial Patterns

As simple examples for relational concepts, we consider relations in which a pairing of entities of the same sort occurs, as in 'older than' or 'larger than'. Gärdenfors (2000) proposes to represent these by constructing a conceptual space with the same qualities occurring on two dimensions.

The same is the case with relations of comparison is ultimately also true for spatial objects that are described by values on the dimensions length, width and height. So let us first turn to the geometrical representation of the property 'x is larger than y', where two dimensions of the type LENGTH occur.[2] As can be seen in figure 4, the resulting graph for this relational concept is simply the half space above the $x = y$ line (see figure 4).

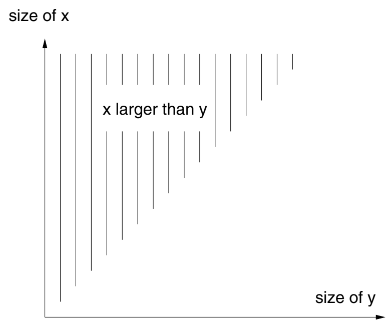

Figure 4: Conceptual space for the 'largerThan(x,y)'-relation

Another example is the relation 'daughterOf(x,y)' in the genealogical tree, which can be represented (in a somehow simplified fashion) as in table 1 (see Gärdenfors, 2000, p. 92ff.). The important point here is only to find a suitable interpretation for the property axes: the dimensions represent the steps backward or forward, respectively, in the genealogy.

(female) successor ↓				common (female) predecessors →
	0	1	2	3
0	(I)	mother	grandmother	great-grandmother
1	daughter	sister	aunt	(grand)aunt
2	granddaughter	niece	cousin	cousin
3	great-granddaughter	grandniece	cousin	cousin

Table 1: Genealogical tree: the 'daughterOf(x,y)'-relation.

[2]This example is also used by Gärdenfors, see p. 92

2.2 Data Reduction is Possible

As already mentioned, the representation of spatial objects or other patterns with height, width, and depth as dimensions is a variant of the same phenomenon that the same qualitative property may occur on different axes. It will certainly be desirable, however, to have a single dimension which encodes, say, the shape of an object. This is indeed possible: a set of data with a multidimensional description can always be translated into a one-dimensional representation if there is a suitable coding mechanism. The human cognitive system is indeed highly effective in coding complex information into simple chunks, and our model would have to reflect this process of creating higher-order properties from simpler ones. In this paper, however, we cannot investigate the details of this process, so let us simply point out that this type of data reduction is always possible in principle: the mathematical fall-back position, as it were, would be *prime number reduction*. This method permits the encoding of values on n dimensions into one single number without any loss of information, using the n first prime numbers. The original information can be recovered via prime factorization:

$$\text{(big) number} = \prod_i p_i^{\text{value on dimension } i}$$

where p_i is the i-th prime number. Example:

$$\text{lengthWidthHeigth} = 2^{\text{length}} \cdot 3^{\text{width}} \cdot 5^{\text{heigth}}.$$

3 Coding of dynamic patterns

Verbs are not only relational concepts but, moreover, they typically denote dynamic situations. So the question arises how this can be captured in the geometric framework. Note that, due to the interpretation of the points in conceptual space, it would not be an option to add a time axis to the coordinate system. Different points in conceptual space always correspond to the mutually exclusive properties of different objects, so the temporal development of one and the same event must be encoded as one single point on a suitable scale, never as an array of several points.

To explain how this can be done, consider the abstract example of a cuboid changing its position, as expressed by the German verbs ‘umfallen’ (to topple over), ‘schwanken’ (to sway) or ‘sich aufrichten’ (to be put upright) (there seems to be a lexical gap in English for an inchoative variant of ‘put upright’).

For one concrete example, let us consider the movement pattern expressed by the sentence:

(1) The cuboid toppled over.

To start, we can imagine a time axis where every point is assigned to a representation of the cuboid in its current position relative to the ground. Formally, we can represent this with values for angles (see figure 5).

Since we have encoded the static representation of the cuboid in a number (the whole cuoid is represented as one point on a scale ⟨cuboid shapes⟩), we now have a two-dimensional representation of the cuboid’s change of position: to each point on the time scale we assign the cuboid’s position relative to the ground (this information can easily be encoded in one number):

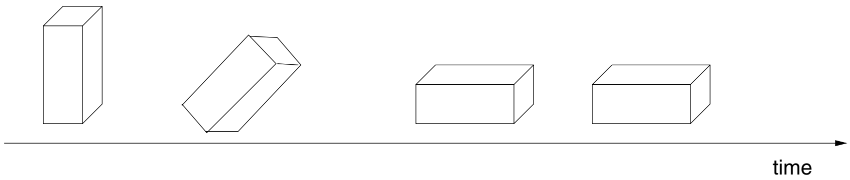

Figure 5: Modelling time in ‘topple over’

t_1	$\cdots$	t_i	t_n	t_{n+1}	$\cdots$
90	$\cdots$	10	0	0	$\cdots$

This array of spatial positions on the time axis can then be treated in the same way as a extended spatial pattern that is stative i. e. it can be encoded as a single property value (whether by prime factorization or by a more psychologically realistic method). Other movement patterns, like swaying, or rising from the ground, can be encoded in an analogous fashion, and a scale for movement types can be constructed (provided that we can posit a suitable ordering for them).

Since this representation only needs complete descriptions of objects, it works in the same way for more complex activities like ‘to walk’. Gärdenfors (p. 98f.) takes up a suggestion from Marr and Vaina, who describe walking as a combination of movement of body parts, where the body parts are modelled as geometrical shapes using the cylinder model of Marr and Nishihara, and the movements are modelled by differential equations.

In sum, then, we have briefly sketched an account of how the geometric model can accommodate dynamic properties. These do not constitute a problem for the framework although they lead to a considerable complexity in the features that are used. In the following sections, we turn to other properties of verb meanings that we think constitute more serious problems for the geometric model.

4 Events and their Properties in Conceptual Space

4.1 The Problem of Commensurability

In this chapter, we explore the applicability of the geometric representation to the descripition and categorisation of events. In doing this, we shall make the simplification of talking about verb meanings, although it is clear that the properties of an event are described by a whole composite expression, consisting of the verb, its arguments, and modifiers. In this paper, we will put aside this compositional aspect for the most part and render verbs in the full semantic potential that is further constrained by the additional material in a sentenc.

If we compare the case of verbs and events to our first simple example of colour terms above, we can immediately note a number of differences that pose problems for the application of the geometric model. For one thing, the semantic features that verb meanings bring along are considerably richer and more abstract than those associated with colours. Verb meanings are more indirectly related to categories of sensory experience. Hence, the quality space is more difficult to determine. A rich set of features may lead to conceptual spaces of high dimensionality; but moreover, it may give rise to a number of different options for classifying verbs into groups. This state of affairs may cause practical as well as theoretical problems.

The major theoretical problem is the following: As already mentioned, a major feature of the geometric approach is that it presents us with a quality space that is divided up among a group of concepts. It is fairly useless to define a conceptual space for the representation of a single verb meaning. Therefore, our analysis of verb meanings must meet two requirements: We have to identify the set of dimensions that define a single verb meaning, and, at the same time, we have to make sure that all those verbs we are comparing share the very same conceptual dimensions. The crucial property of conceptual spaces is that they make available a measure for similarity via the distance between points, and, for geometric reasons, this distance measure would break down if the two verb meanings that we want to compare differ in dimensionality. Figure 6 shows an abstract example of this problem:

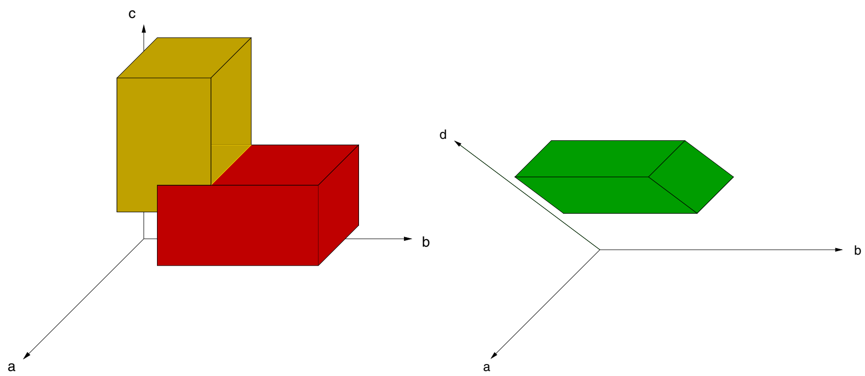

Figure 6: Comparability only in shared dimensions

Assume that verb *A* is described by the property dimensions *a*, *b*, and *c*, and another verb *B* by the dimensions *a*, *b*, and *d*. If the situation is simply this, then no vector is defined from any of the points within *A* to any of those within *B*, and hence no similarity can be ascertained. Without any similarity relation, these two verbs could not belong to the same prototypical category. In the next subsection we shall investigate a concrete set of verbs in order to see whether there is a real danger that verbal categories could indeed fall apart in this way because of the requirements of the model.

Before getting to that, however, we have to consider the theoretical possibility, provided by our geometrical framework, of projecting the two graphs onto the dimensions they have in common. Figure 7 figure illustrates the notion of projection:

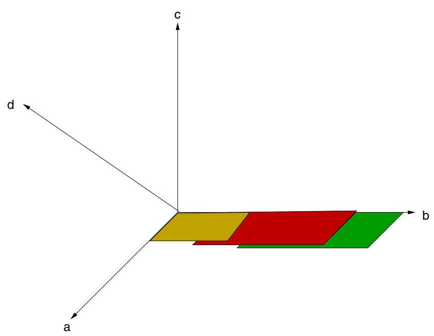

Figure 7: Four axes projected onto two

The similarity measure would be defined for the 2-dimensional projection of the two incommensurable 3D-graphs. Notionally speaking, projection seems to correspond to a licence for ignoring any property dimensions that are not shared between the elements of copmparison, and this would make sense intuitively: this may be the point to bring the fact into play that comparison is always context-dependent. So, setting a context for a comparison—by saying that the objects are similar **with respect to** certain features—might be directly represented by projecting the graphs onto a certain subspace.

However, we have to pay attention to the significance that coordinates in conceptual space bear: they express correlations of properties and as such represent objects. Different points thus stand for different objects, and if this is so, the operation of projection would have to be interpreted as a switch to just a different set of objects, which would seem to be illicit. In the figure above, the projection is onto the plane that has the coordinates $c = d = 0$. So this is the graph of all objects that happen to have the zero value in the dimensions c and d. Even if there are no such objects in the actual world that we could mess up with here, this would have to count as a coincidence. For the geometric representation at least, these hypothetical objects are well-defined. So projection as in figure 7 would not be a way to go. The only thing that could be done is to remove completely all those dimensions that are undefined for any of the objects (see figure 8):

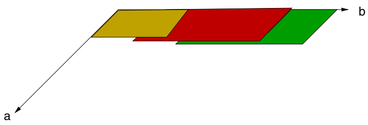

Figure 8: Only two dimensions left

Here, the property dimensions c and d are not present at all, and we are not assuming a zero value in any of these dimensions. In essence, however, what we have done is to construct a new conceptual space. There is no other way of making a comparison than that, in our hypothetical example. Of course, the relevance of these theoretical considerations must be checked with concrete examples. We have seen so far that we have to be careful with the interpretation of the origo of our coordinate system: zero is just another value in a property dimension, it does not express the absence of a feature or the status of its being undefined.

To conclude this section on the geometry of conceptual spaces, it could be seen that the constraints on conceptual representation imposed by the very architecture of the geometric model are non-trivial. Even quite simple considerations indicate that conceptual spaces may involve construction and reconstruction processes if families of concepts come into play whose dimensionality is not uniform for all members (so the model worked much more smoothly for the concept of colour from section 1). However, our discussion so far has only been in the hypothetical mood. In the next section we are going to show that the same general concern will indeed materialise again in the concrete examples for the semantics of verbs.

4.2 A Case Study: "Verbs of Vertical Movement"

In this section we report on a preliminary investigation of a group of German verbs that we have provisionally called "verbs of vertical movement", although ultimately this categorisation will be called into question. This investigation is inspired by Jackendoff's discussion of the English verb 'climb' (see Jackendoff, 1985). The point in Jackendoff's discussion was that the different senses of 'climb' were found to display a family resemblance structure, without any semantic feature that is common to all variants. We shall come back to this point later. To start, let us develop the feature composition of the German translational equivalents of 'climb' by comparing the different variants in a stepwise fashion. First of all, we want to distinguish 'steigen' from 'klettern':

(2) a. auf den Baum klettern.
(on the tree climb)
'to climb (onto) the tree'
b. auf den Berg steigen.
(on the mountain climb)
'to climb (onto) the mountain'
c. die Treppe hinauf steigen.
(the stairs up climb)
'to climb (up) the stairs'
d. der Zeppelin steigt.
(the blimp climbs)
'the blimp is climbing'

For constructing the relevant feature dimensions, we note that the verb 'klettern' (whose meaning seems to be in between the English verbs 'climb' and 'clamber') involves more and more laborious activity. In particular, 'klettern' seems to involve the use of hands and feet while 'steigen' can be used if the situation only involves walking upward (like a hiker's walking up a mountain). It is not immediately clear whether we should use 'effort' or the use of one's extremities as the relevant feature dimension, or both.

$$\begin{bmatrix} \text{DIRECTION} & [\,\textit{upward}] \\ \text{EFFORT} & \langle [\textit{passive}] \ldots [\textit{self-propelled movement}] \ldots [\textit{effortful movement}] \rangle \end{bmatrix}$$

$$\begin{bmatrix} \text{DIRECTION} & [\,\textit{upward}\,] \\ \text{MANNER} & \langle [\textit{null}] \ldots [\textit{feet}] \ldots [\textit{hands\&feet}] \rangle \end{bmatrix}$$

The fact that there is a direct correlation between 'manner' and 'effort' seems to advise against using both—unless redundancy should turn out to be useful for the representation. For the time being, let us content ourselves with a two-dimensional coordinate system (see figure 9):

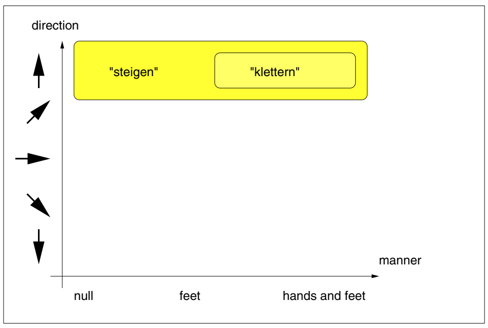

Figure 9: 'klettern' (I)

Here, 'klettern' appears as a sub-region within the representation of the more general verb of upward movement, 'steigen'. To learn more, we have to adduce additional uses that involve directional PPs:

(3) a. Über die Mauer steigen / klettern.
(over the wall climb)

'to climb over the wall'

b. die Leiter hinunter steigen / klettern.
(the ladder down climb)
'to climb down the ladder'

c. an der Regenrinne entlang klettern.
(at the gutter along climb)
'to climb down along the gutter'

d. Die Affen klettern von den Bäumen.
(the monkeys climb from the trees)
'the monkeys are climbing from the trees'

e. *der Ballon steigt nach unten.
(the balloon climbs to down)
'the balloon is climbing down'

These examples show that directions other than 'upward' must be allowed, too. This is clearly conditioned by the PPs, but we take it that we have to represent the verb's potential of being combined with such downward paths as part of the verb meaning. The addition of a PP can then be considered as intersective modification, which carves out a subregion of the verb's extension. The possible appearance of downward PPs plays an important role in Jackendoff's (1985) discussion. We may note in passing that Jackendoff's analysis of these variants as extensions of a prototype category would not absolutely enforce the very same extension on the corresponding German concept. With Jackendoff, we want to emphasise that not all combinations of verbs and downward directionals are possible, as evidenced by the last example.

We now need the following additional values of the quality dimension 'direction':

$$\begin{bmatrix} \text{DIRECTION} \left\langle [\textit{downward}] \ldots [\textit{mixed}] \ldots [\textit{upward}] \right\rangle \\ \ldots \end{bmatrix}$$

The ordering that we impose on this quality dimension is actually not so much due to the direction of the movements as such but to the sum vector that they yield. In this way, purely horizontal movement is treated on a par with truly "mixed" paths of upward and downward movement (however, for brevity we refer to all these cases by the feature 'mixed'). In this way, we can maintain an unambiguous one-dimensional ordering. Note, however, that in admitting these feature values, we can no longer speak of the verbs under investigation as "verbs of vertical movement".

The following graphs (figure 10) show the updated extensions of the verbs 'klettern' and 'steigen', and additionally the subregion that represents the modified cases 'climb down'.

Finally, given that we are now dealing with instances of downward motion, too, we can proceed to include other verbs that denote downward movement, namely the German equivalents (and cognates) of English 'fall' and 'sink'. Obviously, these can be expected to be located in the same conceptual space.

(4) a. Mir ist eine Tasse hinuntergefallen.
(me$_{\text{DAT}}$ has a cup fallen-down)
'I (inadvertently) dropped a cup'

b. Das Schiff sinkt.
(the ship sinks)
'The ship is sinking'

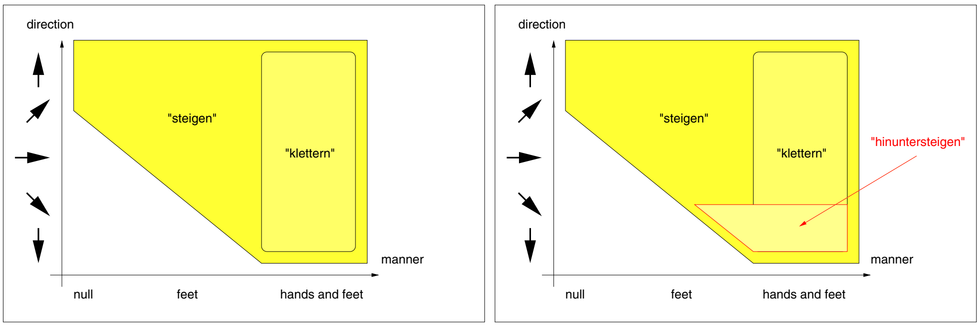

Figure 10: ‘klettern’ (II, III)

c. Der Zucker sinkt durch den Milchschaum.
(the sugar sinks through the milk foam)
‘The sugar is sinking through the milk foam’

d. *Das U-Boot sank an die Oberfläche.
(the submarine sank to the surface)

These verbs must be assigned the following feature composition (including here the possibly redundant dimensions of both manner and effort):

$$\begin{bmatrix} \text{DIRECTION} & [\textit{downward}] \\ \text{MANNER} & [\textit{null}] \\ \text{EFFORT} & [\textit{passive}] \end{bmatrix}$$

They fill the region that has remained empty so far, and we seem to have arrived at a fairly complete partition of the conceptual space of vertical movement.

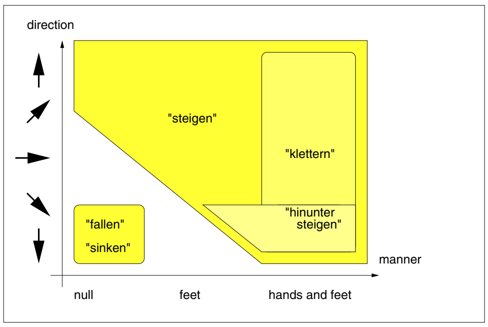

Figure 11: ‘klettern’ (IV)

However, something that is still missing is the distinction between ‘fall’ and ‘sink’. To this end, consider the following data:

(5) a. Das U-Boot sinkt tiefer.
(the submarine is sinking deeper)

b. *Das U-Boot fällt tiefer.
(the submarine is falling deeper)

c. *Das U-Boot klettert hinunter.
(the submarine is climbing down)

d. Das Flugzeug sank tiefer.
(the plane sank deeper)
e. Der aufgewirbelte Staub sank zu Boden.
(the raised dust sank to floor)

We conclude that the difference between these two verbs consists in the way in which the movement interacts with the surrounding medium: sinking is a type of movement that goes through a rather dense medium which slows it down, or it is at least a movement that is slow in comparison to falling (i. e., behaves "as if" it were slowed down, even in empty space). Given this, we might need a new quality dimension like (the viscosity of the) MEDIUM. This in turn presupposes that an object is floating freely, which in the end creates another difference between 'sink' and 'climb down'. Let us provisionally posit this as another dimension. Then, the difference we were looking for is the following: falling is confined to a medium that is gaseous at most, while sinking allows all kinds of medium, provided that the speed is not excessive. For simplicity, we ignore the speed feature in the following:

$$\begin{bmatrix} \pm\text{ATTACHMENT} & [-\ \textit{(floating)}] \\ \text{MEDIUM} & \langle [\textit{null}] \ldots [\textit{thin}] \ldots [\textit{fluid}] \ldots [\textit{viscous}] \rangle \\ \textbf{sinken} & \end{bmatrix}$$

vs.

$$\begin{bmatrix} \pm\text{ATTACHMENT} & [-\ \textit{(floating)}] \\ \text{MEDIUM} & \langle [\textit{null}] \ldots [\textit{thin}] \rangle \\ \textbf{fallen} & \end{bmatrix}$$

Now, there is still a problem: How should we conceive of the interaction of these new quality dimensions with our old verbs 'steigen' and 'klettern', and even of the connection between these two new dimensions themselves?

In the above feature structure, we have (provisionally) employed 'attachment' and 'medium' as two independent dimensions. Let us pursue this assumption somewhat farther. We would now have to assign two new values to verbs like 'klettern' (clamber): the medium in which it proceeds and its mode of attachment. The medium obviously doesn't make any difference, so, for example, you use the same verb whether you climb down a mountain or a submarine rock (we are not aware of any language that makes a distinction here). For 'attachment' we now get the positive value, i. e. the opposite of being freely suspended. Let us adopt the shorthand Ω for the case that a verb permits any value of a feature dimension. Collecting all features, we can then write:

$$\begin{bmatrix} \pm\text{ATTACHMENT} & - \\ \text{MEDIUM} & \langle [\textit{null}] \ldots [\textit{thin}] \rangle \\ \text{DIRECTION} & \downarrow \\ \text{MANNER} & [\textit{null}] \\ \textbf{fallen} & \end{bmatrix} \begin{bmatrix} \pm\text{ATTACHMENT} & + \\ \text{MEDIUM} & \Omega \\ \text{DIRECTION} & \Omega \\ \text{MANNER} & [\textit{hands\&feet}] \\ \textbf{klettern} & \end{bmatrix}$$

However, it is questionable whether the specification of the medium is truly independent of the attachment feature. As already said, we are not aware of any case in which these features vary independently. Rather, it seems that a verb is always either about the grip on a solid object (like 'clamber') or about a relation to a medium. Hence, we could be dealing with different values on one and the same dimension, which therefore exclude each other. So one might suspect that all these verbs in fact describe the way in which a moving entity interacts with its surroundings. You could then either act on a solid "environment", i. e. have contact with an object ('+attachment') or act on a less solid environment, i. e. on a more or less dense medium. In brief: the relevant dimension would be the way in which the moving object exerts force on its environment. We could posit the following scale:

$$\begin{bmatrix} \text{MANNER OF FORCE EXERTION} \left\langle [\textit{null}] \ldots [\textit{buoyancy}] \ldots [\textit{active push-off}] \right\rangle \\ \ldots \end{bmatrix}$$

The verb 'fallen' would be assigned the feature 'null' on this scale, 'sinken' and certain uses of 'steigen' (climb) would get 'buoyancy' and the verb 'klettern' (clamber), along with the other, more active uses of 'steigen' would get 'active push-off'.

The decision between the first account (two independent dimensions) and the second one is not easy, although it appears that the second one has something to recommend itself. So we may notice how little restrictions there are on the ways of constructing property dimensions; the arguments that can be adduced to decide the point will often consist of considerations like the economy of the representation. And actually, examining things even closer, we can find arguments that rather support a third view on the matter just discussed. Thus, we believe that positing a dimension like MANNER OF FORCE EXERTION does not account for certain correlations. Note the following: Those verbs that allow the DIRECTION feature to vary freely are exactly those that involve what we have just called 'active push-off'. Conversely, those verbs that are sensitive to the viscosity of a medium and describe movement of free-floating objects always carry a fixed specification for direction.

Furthermore, there is a difference in the significance of the environmental object/medium that is not captured by positing a single unified quality dimension of force exertion: For the verbs which are '+attachment', the solid object functions as a reference object. The direction of the path is largely determined by the orientation of this object. However, the path denoted by verbs like 'fall' is inherent in the verb meaning and not related to the medium (which makes sense because the medium itself does not give rise to a designated direction). Our ultimate conclusion is therefore that there are two distinct groups of verbs: verbs of "push-off" with respect to a reference object (like 'klettern') and verbs of inherent vertical motion, which can be sensitive to properties of the medium (like 'sinken'). The verb 'steigen' as well as 'climb' must now be considered as polysemous: they have both a "push-off" reading and a "vertical motion" reading, and depending on the reading, they acquire all properties of the respective verb class (including the requirement that "vertical motion" verbs are non-agentive, for example). With this finding, we contradict Jackendoff (1985) who assumed that there is a family-resemblance structure. The absence of readings with downward direction of 'climb' if it denotes passive movement (i. e., buoyancy) was explained by Jackendoff (1985) by the assumption that this is required by the family resemblance structure of a prototypical category; so, a reading that involved both buoyancy and downward movement would not bear a sufficiently strong similarity relation to the prototype any more. However, this position does not seem to explain the similarity between the functioning of each reading of 'climb' and other words of the same relevant class that we have discovered. The polysemy view, we think, is additionally made plausible by the fact that it seems impossible to predict that the German verb 'steigen' in the "push-off" sense is focused on movement supported by the feet but not by necessarily on support by the hands, in distinction to 'klettern'. This is an idiosyncrasy of the "push-off" sense and is in line with the fact that these verbs, but not those of vertical motion, tend to make specifications as to the manner of movement.

We can sum up our result in the following table (where we additionally include 'rutschen' (slide)):

Group 1: Movement; contact with an object

	"klettern"	"steigen"	"rutschen"
EFFORT	'maximal'	'medium'	0
MANNER	'Hands & Feet'	'Feet'	'sliding'
DIRECTION	Ω	Ω	Ω

Group 2: Movement; in a medium

	"steigen"	"sinken"	"fallen"
MEDIUM	$\langle$empty... viscous$\rangle$	$\langle$empty... viscous$\rangle$	$\langle$empty... thin$\rangle$
ENERGY	0+	0+	0
DIRECTION	$\uparrow$	$\downarrow$	$\downarrow$

Translating this into the geometrical representation, we now have to posit two distinct conceptual spaces. Calling the verb 'climb' polysemous means that it occurs in both, but plays entirely different roles in the partitioning of the respective conceptual space:

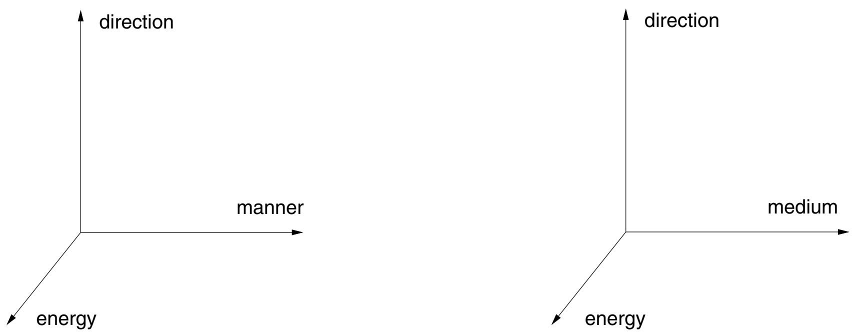

Figure 12: Distinct conceptual spaces

5 Conclusion

Putting together the results from the theoretical considerations and the empirical investigations of verb meanings, we find that they converge on the same general point. Conceptual spaces are designed to represent homogeneous semantic fields that are partitioned by a number of concepts (and, consequently, linguistic items). So it is inherent in the model that it describes whole semantic fields, not individual lexical items. Higher-order concepts can exhibit more variation in their featural makeup, and as soon as the model is applied to such cases, the dimensional homogeneity that is required in the architecture of the geometric model may be lost. Verbs, in particular, can be found to pose intricate problems for this kind of representation because, first of all, it is not sure what kind of quality dimensions will be called for. At present, there seem to be no restrictions as to what can count as a "natural" property in which to define a verb's conceptual space. Exploring the question of which dimensions make up the conceptual representation of verbs like ‘climb’, we ultimately took our guidance from heuristic considerations such as the ability to capture certain generalisations in an efficient way.

We have argued that the optimal analysis of verbs like ‘climb’ lies in recognising two distinct semantic fields, or semantic spaces, and this result brings us right back to our initial considerations of the architecture of the geometrical representation. If it is true that we end up with two distinct conceptual spaces, we still need a way of stating similarities like the downward directionality shared by, say, ‘sink’ and ‘climb down’. This state of affairs seems to constitute one example of the case that we had first delineated in an abstract way, namely that similarity statements can often not be read off from the distance between points in a unified conceptual space, because the spaces in question are incommensurable. So additional operations are required that blend distinct spaces into unified ones of lower dimensionality. The latter can serve as a base for similarity judgements, but not as a lexical-conceptual representation, due to their strongly reduced content.

As another observation to a similar effect, it is obvious that one and the same property can reappear in different conceptual spaces (like agentivity or, again, “direction of movement”), and this seems to reflect the fact that those properties are the really fundamental ingredients of semantic representation while conceptual spaces are derived entities. In fact, Gädenfors (2000) himself points out that higher-order concepts can be constructed on the basis of more fundamental ones (ultimately on the basis of elementary sensory and intentional ones with a fixed and hard-wired dimensional makeup). So, all in all, the aspect of active construction of conceptual spaces must be strongly emphasised.

Employing conceptual spaces in order to give a lexical analysis, as we tried to do for verbs like ‘climb’, actually has to be viewed as nothing more than an attempt to construct a particular conceptual space that optimally serves a particular purpose. What the “lexical representation” did was to try to establish a particular grouping of items (by way of a particular selection of feature dimensions) so as to maximise comparability—at the expense of broadness of coverage. But maybe this is a deeper point about lexical representations, which is not to blame on the geometric model: that conceptual features do not exist in isolation but can only be established in the course of comparison operations.

References

[Gärdenfors, 2000] Gärdenfors, P. (2000). *Conceptual Spaces: the geometry of thought*. MIT press, Cambridge Mass.

[Hard and Sivik, 1981] Hard, A. and Sivik, L. (1981). NCS–natural color system: a Swedish standard for color notation. *Color Research and Application*, 6:129–138.

[Jackendoff, 1985] Jackendoff, R. S. (1985). Multiple Subcategorization and the Θ-Criterion: The Case of CLIMB. *Natural Language and Linguistic Theory*, 3:271–295.

[L. Sivik and C. Taft, 1994] L. Sivik and C. Taft (1994). Color naming: a mapping in the NCS of common color terms. *Scandinavian Journal of Psychology*, 35:144–164.

[Tversky, 1977] Tversky, A. (1977). Features of Similarity. *Psychological Review*, 8:327–352.