Semantics Based on Conceptual Spaces

Peter Gärdenfors

Lund University Cognitive Science, Kungshuset, S-22222 Lund, Sweden
`Peter.Gardenfors@lucs.lu.se`

Abstract. The overall goal is to show that conceptual spaces are more promising than other ways of modelling the semantics of natural language. In particular, I will show how they can be used to model actions and events. I will also outline how conceptual spaces provide a cognitive grounding for *word classes*, including nouns, adjectives, prepositions and verbs.

1 Introduction

Within traditional philosophy of language, semantics is seen as mapping between language and the world (or several "possible worlds"). This view has severe problems. For one thing, it does not involve the users of the language. In particular, it does not tell us anything about how individual users can grasp the meanings determined by such a mapping (Harnad 1990, Gärdenfors 1997).

Another tradition, *cognitive semantics*, brings in the language user by focusing on the relations between linguistic expressions and the user's mental representation of their meanings. According to cognitive semantics, the meanings of words are represented specifically as *image schemas*. These schemas are abstract mental pictures with an inherent spatial structure, constructed from elementary topological and geometrical structures like "container," "link", and "source-path-goal." Such schemas are commonly assumed to constitute the representational form common to perception, memory, and semantic meaning.

Although there have been some attempts to construct computational models of image schemas (e.g. Holmqvist 1993), they are not well suited for formal modelling. In particular, they are not well developed for handling dynamic entities, such as actions and events. In this article, I will model actions and events using conceptual spaces (Gärdenfors 2000). My goal is to show that conceptual spaces show more promise than other ways of modeling the semantics of natural language (see also Gärdenfors (1996)). I will further show how they can provide a cognitive grounding for *word classes*. In linguistics, word classes are defined by syntactic criteria. However, a theory of cognitive semantics worthy of its name should at least be able to explain the main categories of words – i.e., nouns, adjectives, prepositions, and verbs -- in terms of cognitive mechanisms. I will outline such an account.

2 Conceptual Spaces as a Semantic Framework

A given conceptual space consists of a number of *quality dimensions*. Examples of quality dimensions are temperature, weight, brightness, pitch, and force, as well as the

M. Banerjee and A. Seth (Eds.): ICLA 2011, LNAI 6521, pp. 1–11, 2011.

three ordinary spatial dimensions of height, width, and depth. Some quality dimensions are of an abstract non-sensory character. One aim of this article is to argue that force dimensions are essential for the analysis of actions and events.

Quality dimensions correspond to the different ways stimuli can be judged similar or different. For example, one can judge tones by their pitch, and that will generate a certain ordering of the auditory perceptions. As a general assumption, the smaller the distance between the representations of two objects, the more similar they are. The coordinates of a point within a conceptual space represent particular instances along each dimension: for example, a particular temperature, a particular weight, etc. For simplicity, I assume that the dimensions have some metric, so that one can talk about distances in the conceptual space. Such distances indicate degrees of similarity between the objects represented in the space.

It is further assumed that each of the quality dimensions can be described in terms of certain geometrical shapes. A psychologically interesting example is colour. Our cognitive representation of colour can be described along three dimensions. The first is hue, represented by the familiar colour circle going from red to yellow to green to blue, then back to red again. The topology of this dimension is thus different from the quality dimensions representing time or weight, which are isomorphic to the real number line. The second dimension is saturation, which ranges from grey at the one extreme, to increasingly greater intensities of colour at the other. This dimension is isomorphic to an interval of the real number line. The third dimension is brightness, which varies from white to black, and thus is also isomorphic to a bounded interval of the real number line. Together, these three dimensions---one circular, two linear---constitute the colour domain as a subspace of our perceptual conceptual space. It is typically illustrated by the so-called colour spindle.

The primary function of the dimensions is to represent various qualities of objects in different domains. Since the notion of a domain is central to the analysis, I should give it a more precise meaning. To do this, I will rely on the notions of separable and integral dimensions, which I take from cognitive psychology (Garner 1974, Maddox 1992, Melara 1992). Certain quality dimensions are *integral*: one cannot assign an object a value on one dimension without giving it a value on the other(s). For example, an object cannot be given a hue without also giving it a brightness (and a saturation). Likewise the pitch of a sound always goes with a particular loudness. Dimensions that are not integral are *separable*: for example, the size and hue dimensions. Using this distinction, a *domain* can now be defined as a set of integral dimensions that are separable from all other dimensions.

A *conceptual space* can then be defined as a collection of quality dimensions divided into domains. However, the dimensions of a conceptual space should not be seen as fully independent entities. Rather, they are correlated in various ways, since the properties of those objects modelled in the space co-vary. For example, in the fruit domain, the ripeness and colour dimensions co-vary.

It is impossible to provide any complete listing of the quality dimensions involved in the conceptual spaces of humans. Learning new concepts often means expanding one's conceptual space with new quality dimensions (Smith 1989).

3 Properties and Concepts

Conceptual spaces theory will next be used to define a *property*. The following criterion was proposed in Gärdenfors (1990, 2000), where the geometrical characteristics of the quality dimensions are used to introduce a spatial structure to properties:

Criterion P: A *natural property* is a convex region in some domain.

The motivation for Criterion *P* is that, if some objects located at *x* and *y* in relation to some quality dimension(s) are both examples of a concept, then any object that is located between *x* and *y* with respect to the same quality dimension(s) will also be an example of the concept.

Properties, as defined by criterion *P*, form a special case of *concepts*. I define this distinction in Gärdenfors (2000) by saying that a property is based on a single domain, while a concept is based on one *or more* domains. This distinction has been obliterated in both symbolic and connectionist accounts, which have dominated the discussions in cognitive science. So for example, both properties and concepts are represented by predicates in first-order logic. However, the predicates of first-order logic correspond to several different grammatical categories in natural language, most importantly those of adjectives, nouns, and verbs.

A paradigm example of a concept that is represented in several domains is "apple" (compare Smith et al. 1988). One of the first problems when representing a concept is to decide which are the relevant domains. When we encounter apples as children, the first domains we learn about are, presumably, those of colour, shape, texture, and taste. Later, we learn about apples as fruits (biology), about apples as things with nutritional value, etc.

The next problem is to determine the geometric structure of the domains: i.e., which are the relevant quality dimensions. Taste space can be represented by the four dimensions of sweet, sour, salty, and bitter; the colour domain by hue, saturation, and brightness. Other domains are trickier. For example, it is difficult to say much about the topological structure of "fruit space", in part because fruits (such as apples) can be described relative to several domains. Some ideas about how such "shape spaces" should be modelled have been discussed in e.g. Marr and Nishihara (1978), Edelman (1999), and Gärdenfors (2000). Instead of offering a detailed image of the structures of the different domains, let me represent the "apple" regions in the domains verbally, as follows:

Domain	*Region*
colour	red-yellow-orange
shape	roundish
texture	smooth
taste	regions of the sweet and sour dimensions
nutrition	values of sugar content, fibre content, vitamins, etc.
fruit	specification of seed structure, fleshiness, peel type, etc.

Concepts are not just bundles of properties. They are also *correlations* between regions from different domains that are associated with the concept. The "apple" concept has a strong positive correlation between sweetness in the taste domain and sugar content in the nutrition domain, and a weaker positive correlation between redness and sweetness.

Such considerations motivate the following definition for concepts. (For a more precise definition, see Chapter 4 in Gärdenfors 2000.)

Criterion C: A *concept* is represented as a set of convex regions in a number of domains, together with information about how the regions in different domains are correlated.

Elements from theories in psychology and linguistics contribute to the analysis of concepts I present here. The kind of representation intended by Criterion *C* is, on the surface, similar to *frames*, with slots for different *features* (sometimes called slots, attributes, or roles; see for example Noy and McGuinness (2001)). Frames have been very popular within cognitive science as well as in linguistics and computer science. However, Criterion *C* is richer than frames, since it allows representing concepts as more or less *similar* to each other and objects (instances) as more or less representative of a concept. Conceptual spaces theory can be seen as combining frame theory with prototype theory, although the geometry of the domains makes possible inferences that cannot be made in either of those theories (Gärdenfors 1990, 2000).

4 The Semantics of Adjectives, Nouns, and Prepositions

Next I will outline how analysing properties and concepts in terms of conceptual spaces can provide a *cognitive grounding* for different word classes. In this section I discuss adjectives, nouns, and prepositions. I will discuss verbs later.

The main semantic difference between adjectives and nouns is that adjectives (e.g., "red," "tall," "round") normally refer to a single domain and represent properties; while nouns (e.g., "dog," "apple", "town") normally relate to several domains. (Verbs, unlike nouns or adjectives, are characterized by their dynamic content, which I will analyze in terms of actions based on the force domain.)

Most properties expressed by adjectives in natural languages are natural properties according to Criterion *P*. For instance, in Gärdenfors (2000) I conjectured that all colour terms in natural languages express natural properties with respect to the the colour dimensions of hue, saturation, and brightness. This means that there should be no language which has a single word for the colours denoted by "green" and "orange" in English (and which includes no other colours), since such a word would represent two disjoint areas in the colour space. Sivik and Taft (1994) and Jäger (2009) have provided strong support for this conjecture. Their studies follow up on the investigations of basic colour terms by Berlin and Kay (1969), who compared and systematized terms from a wide variety of languages. Jäger (2009) studied colour classification data from 110 languages and found a median value of 93.6% correct classifications in an optimally convex partitioning of the colour space. Given the statistical aberrations in the data, this is a very high figure, which gives strong support to Criterion *P* at least within the domain of colours.

A paradigmatic example of the semantics for nouns is the analysis of “apple” from the previous section. Nouns do not only denote physical objects as located within a limited spatial region: consider, for example, “thunder”, “family”, and “language,” let alone more abstract nouns. A noun typically denotes a phenomenon with *correlations* across a number of domains: in other words, nouns are represented by clusters in the conceptual space. Not all potential such clusters will be named by nouns in a language; an important factor is whether the correlations are pragmatically significant: that is, whether they are helpful in choosing the right actions.

Prepositions have likewise been studied extensively within cognitive semantics (for example, Herskovits 1986, Lakoff 1987, Landau and Jackendoff 1993, Zwarts 1995, Zwarts and Winter 2000, Zwarts to appear). A locative preposition (e.g., “in front of”) combines with a noun phrase (e.g., “the castle”) that refers to a spatially located object. The preposition maps the reference object to a region that is related to the object. (This criterion is put forward by e.g. Jackendoff 1983 and Landau and Jackendoff 1993, p. 223). Zwarts (1995) proposes to analyse this region as a set of vectors radiating from the reference object.

The basic semantic function of prepositions is to express spatial relations; but they are also used in a number of metaphorical and metonymic ways. Landau and Jackendoff (1993) offer a neuro-linguistic explanation. They propose two distinct cognitive systems: one for *objects* (the “what” system), and one for *places* (the “where” system). These systems relate to two different pathways in the visual cortex. The separation of the systems results in the separation between the *nominal* and *prepositional* systems in language. Zwarts (to appear) argues that, in some contexts, the force domain is also necessary to analyse the meaning of prepositions, meaning that their semantics cannot be handled solely as spatial relations. To what extent one can find neuro-scientific support for the representation of force remains to be seen.

5 Modelling Actions

One idea for a model of actions comes from Marr and Vaina (1982) and is explored further in Vaina (1983). Marr and Vaina extend Marr and Nishihara’s (1978) cylinder models of objects to an analysis of actions. In Marr and Vaina’s model, an action – say, a person walking – is described via differential equations for the movements of the implicated body parts.

It is clear that these equations can be derived, by application of Newtonian mechanics, from the forces that are applied to the legs, arms, and other moving body parts. Our cognitive apparatus is not precisely built for thinking in terms of Newtonian mechanics, but I hypothesize that, nevertheless, our brains successfully extract the forces that lie behind different kinds of movement-involving action. I will present some support for this hypothesis below. More precisely, I submit, building on Gärdenfors (2007), that the fundamental cognitive representation of any action consists of the *pattern of forces* that generates it. I speak of patterns of forces, since, for bodily motions, several body parts are involved; and thus, several force vectors are interacting (in analogy with Marr and Vaina’s differential equations). It should be emphasized, however, that the “forces” represented by the brain are psychological and not the scientific construct introduced by Newton.

The patterns of forces can be represented in principally the same way as the patterns of shapes discussed earlier. For example, the pattern of force of a person running is different from the pattern of a person walking; and the pattern for saluting is different from that of throwing (Vaina and Bennour 1985).

The best source of empirical support for my hypothesis comes from psychophysics. During the 1950's, Gunnar Johansson developed his patch-light technique for analyzing biological motion without any direct information about shape. (For a survey, see Johansson 1973.) He attached light bulbs to the joints of actors who were dressed in black and moved in a black room. The actors were filmed performing various actions, such as walking, running, and dancing. Subjects watching the films, in which only the dots of light could be seen, recognized the actions within tenths of a second. Furthermore, the movements of the dots were immediately interpreted as coming from the actions of a human being. Further experiments by Runesson and Frykholm (1981, 1983) showed that subjects were able to extract subtle details about the actions, such as the gender of walkers or the weight of lifted objects (where the objects, like the actors themselves, were not visible).

One lesson to be learned from the experiments by Johansson and his followers is that the kinematics of a movement contains sufficient information for identifying the underlying dynamic patterns of force. Runesson (1994, pp. 386-387) claims we can directly perceive the forces that control various kinds of motion. He calls this principle the *kinematic specification of dynamics*, according to which the kinematics of a movement contains sufficient information to identify the underlying dynamic patterns of force. It is obvious that his principle accords well with the representation of actions that is proposed here. (Note however that Runesson takes a Gibsonian perspective on the perceptual information available, which means he would find it methodologically unnecessary to consider such mental constructions as conceptual spaces.)

Even though the empirical evidence is incomplete, my proposal is that, by adding force dimensions to a conceptual space, one obtains the basic tools for analyzing the dynamic properties of actions. The forces involved need not only be physical forces, but also *emotional* or *social* forces.

To identify the structure of the action space, one should investigate similarities between actions. This can be done with basically the same methods as for investigating similarities between objects: e.g., "walking" is more similar to "running" than to "throwing". Little is known about the geometrical structure of the action space. I make the weak assumption that the notion of betweenness remains meaningful. This allows me to formulate the following criterion, in analogy with Criterion *C*:

Criterion A: An *action category* is represented as a convex region in action space.

Of course, the more forces are involved, the more multi-dimensional the action space will be and the more complicated it will be to identify the relevant convex regions. One way to support the connection between Criterion *C* and Criterion *A* is to establish that action categories share a similar structure with object categories, as Hemeren (2008, p. 25) has suggested. In a series of experiments (1996, 1997, 2008), he showed that action categories have a similar hierarchical structure and show similar typicality effects to object concepts. Overall, there are strong reasons to believe that actions exhibit many of the *prototype effects* that Rosch (1975) described for object categories.

6 A Two-Vector Model of Events

In keeping with Gärdenfors and Warglien (submitted), I want briefly to present a model of events that is likewise based on conceptual spaces. Events are treated as complex structures that build on conceptual spaces, in particular the action space. The starting point is that all events involve an *agent* and a *patient*.

Agents and patients are modelled as (material or non-material) objects, and can therefore be represented as points in conceptual spaces. The domains of the spaces determine the relevant properties of the agent and the patient. An agent is an animate or inanimate object. Even though I am not providing any analysis of causation here, the common understanding is that the agent is the one causing something to happen. (of course, one should allow that the action can be null, in the case of an event that is a state). An event is individuated by the further understanding that the agent causes the event to happen independently of other events.

An agent is described by an *agent space* that at minimum contains a force domain in which the action performed by the agent can be represented (this is the assumption of *agency*). Following the analysis from the previous section, I will model an action as a force vector (or, more particularly, as a pattern of forces). The agent space may also contain a physical space domain that assigns the agent a location. In particular, in the special case when patient = agent -- i.e., the agent is doing something to itself -- the properties of the agent must be modelled.

A *patient* is again an animate or inanimate object. The patient is described by a *patient space* that contains the domains needed to account for those properties of the patient relevant to the event that is modelled. The properties often include the location of the patient and sometimes its emotional state. A force vector is associated with the patient and represents the (counter-)force exerted by the patient in relation to the agent's action. This can be an intentionless physical force, as when a door does not open when pushed; or it can be an intentionally generated force, as when a person pushes back upon being pushed. For many events, the representation of the patient's force vector is unknown and can be left unspecified; or else it can be taken as prototypical, entailing that the consequences of the agent's action are left open to various degrees.

The force exerted by the agent's action will change one or more properties of the patient. The elementary operations possible on vectors provide a reasonable account for how changes can result from compositions of forces from the agent and the patient. The resultant force vector is the $r = f + c$, where f is the force generated by the agent's action and c is the counter-force of the patient. We then define an event as a mapping between an action in an agent space and a resulting change in a patient space that is the result of applying r. Central to the event are the changes to properties in other domains of the patient space. For example, the location of the patient may change; or its colour may change, if the event involves the action of painting.

This way of representing things makes an explicit difference between an action that is mapped into the patient space and the force exerted by such action. Two different actions (e.g. kicking and punching) might produce the same force vector r in the patient space. It might not be sufficient for characterizing an event to represent the force composition of f and c; the initiating action from the agent must also be represented. This will become even more relevant as event categories are introduced below.

As a simple example, consider the event of *pushing a table*. In such an event, the agent (a person) applies a physical force to a patient (a table). The result is a change in the location of the patient and thus a change in its properties (unless there are balancing counter-forces present, such that the resulting change vector is zero). The change vector depends on the properties of the patient and other aspects of the surrounding world (for example, friction). Another example is an event of walking. In this case, the agent and the patient are identical, so the agent applies a force to itself.

Some actions are ongoing: the agent exerts the force for an unbounded period of time, for example by walking or pushing an object, with the consequence that there may be no definite end point to the changes in the patient space. This is a special case of a more general type of event: *processes*. In bounded events, the agent's force vector is applied for a limited time period, and the change vector in patient space has a clear end point. In many languages, the difference between processes and bounded events is reflected in the syntax: for example, by various forms of aspect. The focus here is on bounded events, but most of the elements of the event representations will apply to corresponding representations of (unbounded) processes.

In general, events should be represented not only as single spatiotemporally located instances, but also as event categories, like "climbing a mountain". I will next provide a formal framework for analysing event categories, of which single events can be considered as instances.

The earlier description of the change vector can be generalized to that of a *change vector field.* The change vector field associates to each point in the patient space that vector change induced by a particular action, taking into account, if necessary, the (counter-)force exerted by the patient. An event category then represents how the agent space potentially affects the patient vector field. For example, the event category of pushing a table should represent the effect of different, albeit similar, patterns of force on the different points in the table patient space.

Event categories can be represented at different levels: there are subcategories of events just as there are of objects. For example, "pushing a door open" is a subcategory of "pushing a door" where the agent force exceeds the counter-force of the patient. "Pushing but failing to open a door" is another subcategory, one where the counter-force cancels out the agent force.

For many kinds of events where the focus is on the changes in the patient, the identity of the agent can be and often is ignored. For example, in the event of somebody falling ill, the cause of illness is often not considered. Similarly, if the force vector is null (i.e., the event is a state), the identity of the agent is irrelevant.

This model of events and event categories is presented in greater detail in Gärdenfors and Warglien (submitted). What is new here, apart from using conceptual spaces as the general supporting framework, is the introduction of the two vectors as forming the core of an event.

7 The Role of Events in the Semantics of Verbs

The fundamental connection between the semantics of natural language and events is that a simple sentence *typically expresses an event.* For this reason, events are central units in any theory of semantics. For this reason, a typical sentence contains the basic building blocks of subject, object and verb, corresponding to the agent, patient and vectors of my model. A single verb can never completely describe an event, but only bring out one aspect of it. I propose that a verb represents *one of the vectors in the model of an event.* In linguistics, a distinction is often made between *manner* and *result* verbs. I suggest that if the verb focuses on the force vector of the agent, as for example in "push" or "hit", then it is a *manner* verb; while if it focuses on the change vector of the patient, as for example in "move" or "stretch", it is a *result* verb.

In the cognitive semantics tradition of Lakoff (1987) and Langacker (1987), the focus has been on the spatial structure of the image schemas only (the very name suggests this), with no attempt to represent the forces involved in the event. This is an essential departure from the way I have proposed modelling actions and events.

Talmy (1988) presents an alternative model of action and interaction. Talmy emphasizes the role of forces and dynamic patterns in image schemas through what he calls *force dynamics*. He develops a schematic-based formalism that allows him to represent the difference of force patterns in expressions like "the ball kept rolling because of the wind blowing on it" and "the ball kept rolling despite the stiff grass". Interactions between agent (what he calls the *agonist*) and Patient (the *antagonist*) are central to his framework as they are to the one presented here. However, some important differences should be highlighted. First is the role that spaces and mappings between spaces play. While Talmy's force dynamics are situated in generic spaces, I am grounding the semantics of events in a theory of conceptual spaces and of the mappings between them. This creates a more flexible and comprehensive framework, one that can take into account the qualitative dimensions of the agent's actions and the changes in the patient. Second, my framework makes a natural distinction between single events and event categories, and is able to account for those events in which the reaction of the patient is not specified.

Finally, an important advantage of the spatial representation of the event structure I have presented is that it allows one to map from one event type to another, comparing their structure and making it possible to address, for example, the metaphorical use of an action or event. In Warglien and Gärdenfors (submitted), we explore the use of topological tools to model such metaphorical mappings.

8 Conclusion

This paper has been of a programmatic nature, advocating an approach to the semantics of natural languages based on conceptual spaces. I have outlined how properties, concepts, actions, and events can all be modelled and a cognitive semantics for different word classes generated.

The model of events combines the analysis of nouns, in terms of concepts for representing agents and patients, with the analysis of actions. Such an analysis shows the intimate connection between how we cognitively represent actions and objects, One that is reflected in the linguistic tools we use for expressing actions and events. I propose that the force vector produced by the agent and the result vector induced in the patient reflect the distinction between manner and result verbs. More abstractly, events represent what philosophers call propositions: that is, the semantic contents of basic sentences.

I have tried to show how conceptual spaces in general, and their application to the force domain in particular, can be useful tools for sharpening cognitive semantics. With the aid of the topological and geometric structure of the various domains, a better foundation for the concept of image schemas is obtained. This applies in particular to dynamic schemas.

Acknowledgements

I gratefully acknowledge support from the VAAG project under the ESF framework LogiCCC and the Swedish Research Council for support under a Senior Individual Grant and the Linnaeus environment *Thinking in Time: Cognition, Communication and Learning*.

References

1. Berlin, B., Kay, P.: Basic Color Terms: Their Universality and Evolution. University of California Press, Berkeley (1969)
2. Edelman, S.: Representation and Recognition in Vision. MIT Press, Cambridge (1999)
3. Gärdenfors, P.: Induction, conceptual spaces and AI. Philosophy of Science 57, 78–95 (1990)
4. Gärdenfors, P.: Conceptual spaces as a basis for cognitive semantics. In: Clark, A., et al. (eds.) Philosophy and Cognitive Science, pp. 159–180. Kluwer, Dordrecht (1996)
5. Gärdenfors, P.: Does semantics need reality?. In: Does Representation Need Reality? Proceedings of the International Conference New Trends in Cognitive Science, Austrian Society of Cognitive Science Technical Report 97-01, Vienna, pp. 113–120 (1997)
6. Gärdenfors, P.: Conceptual Spaces: The Geometry of Thought. MIT Press, Cambridge (2000)
7. Gärdenfors, P.: Representing actions and functional properties in conceptual spaces. In: Ziemke, T., Zlatev, J., Frank, R.M. (eds.) Body, Language and Mind. Embodiment, vol. 1, pp. 167–195. Mouton de Gruyter, Berlin (2007)
8. Gärdenfors, P., Warglien, M.: Using conceptual spaces to model actions and events. Manuscript (submitted)
9. Garner, W.R.: The Processing of Information and Structure. Erlbaum, Potomac (1974)
10. Harnad, S.: The symbol grounding problem. Physica D 42, 335–346 (1990)
11. Hemeren, P.E.: Frequency, ordinal position and semantic distance as measures of cross-cultural stability and hierarchies for action verbs. Acta Psychologica 91, 39–66 (1996)
12. Hemeren, P.E.: Typicality and context effects in action categories. In: Proceedings of the 19th Annual Conference of the Cognitive Science Society, p. 949. Lawrence Erlbaum Associates, Stanford (1997)
13. Hemeren, P.E.: Mind in Action. Lund University Cognitive Studies 140, Lund (2008)
14. Herskovits, A.: Language and Spatial Cognition: An Interdisciplinary Study of the Prepositions in English. Cambridge University Press, Cambridge (1996)
15. Holmqvist, K.: Implementing Cognitive Semantics. Lund University Cognitive Studies 17, Lund (1993)
16. Jackendoff, R.: Semantics and Cognition. MIT Press, Cambridge (1983)
17. Jäger, G.: Natural colour categories are convex sets. In: Aloni, M., Bastiaanse, H., de Jager, T., Schulz, K. (eds.) Logic, Language and Meaning. LNCS (LNAI), vol. 6042, pp. 11–20. Springer, Heidelberg (2010)

18. Johansson, G.: Visual perception of biological motion and a model for its analysis. Perception and Psychophysics 14, 201–211 (1973)
19. Lakoff, G.: Women, Fire, and Dangerous Things. The University of Chicago Press, Chicago (1987)
20. Landau, B., Jackendoff, R.: "What" and "where" in spatial language and spatial cognition. Behavioral and Brain Sciences 16, 217–238 (1993)
21. Langacker, R.W.: Foundations of Cognitive Grammar. Theoretical Prerequisites, vol. I. Stanford University Press, Stanford (1987)
22. Maddox, W.T.: Perceptual and decisional separability. In: Ashby, G.F. (ed.) Multidimensional Models of Perception and Cognition, pp. 147–180. Lawrence Erlbaum, Hillsdale (1992)
23. Marr, D., Nishihara, H.K.: Representation and recognition of the spatial organization of three-dimensional shapes. Proceedings of the Royal Society in London B 200, 269–294 (1978)
24. Marr, D., Vaina, L.: Representation and recognition of the movements of shapes. Proceedings of the Royal Society in London B 214, 501–524 (1982)
25. Melara, R.D.: The concept of perceptual similarity: from psychophysics to cognitive psychology. In: Algom, D. (ed.) Psychophysical Approaches to Cognition, pp. 303–388. Elsevier, Amsterdam (1992)
26. Noy, N.F., McGuinness, D.L.: Ontology development 101: A guide to creating your first ontology, Stanford Knowledge Systems Laboratory Technical Report, Stanford, CA (2001)
27. Rosch, E.: Cognitive representations of semantic categories. Journal of Experimental Psychology: General 104, 192–233 (1975)
28. Runesson, S.: Perception of biological motion: The KSD-principle and the implications of a distal versus proximal approach. In: Jansson, G., Bergström, S.-S., Epstein, W. (eds.) Perceiving Evens and Objects, pp. 383–405. Lawrence Erlbaum Associates, Hillsdale (1994)
29. Runesson, S., Frykholm, G.: Visual perception of lifted weights. Journal of Experimental Psychology: Human Perception and Performance 7, 733–740 (1981)
30. Runesson, S., Frykholm, G.: Kinematic specification of dynamics as an informational basis for person and action perception. Expectation, gender recognition, and deceptive intention. Journal of Experimental Psychology: General 112, 585–615 (1983)
31. Sivik, L., Taft, C.: Color naming: a mapping in the NCS of common color terms. Scandinavian Journal of Psychology 35, 144–164 (1994)
32. Smith, L.: From global similarities to kinds of similarities – the construction of dimensions in development. In: Vosniadou, S., Ortony, A. (eds.) Similarity and Analogical Reasoning, pp. 146–178. Cambridge University Press, Cambridge (1989)
33. Smith, E.E., Osherson, D.N., Rips, L.J., Keane, M.: Combining prototypes: a selective modification model. Cognitive Science 12, 485–527 (1988)
34. Talmy, L.: Force dynamics in language and cognition. Cognitive Science 12, 49–100 (1988)
35. Vaina, L.: From shapes and movements to objects and actions. Synthese 54, 3–36 (1983)
36. Vaina, L., Bennour, Y.: A computational approach to visual recognition of arm movement. Perceptual and Motor Skills 60, 203–228 (1985)
37. Warglien, M., Gärdenfors, P.: Semantics, conceptual spaces and the meeting of minds. Manuscript (submitted)
38. Zwarts, J.: The semantics of relative position. In: Proceedings of Semantics and Linguistic Theory, vol. 5, pp. 405–422 (1995)
39. Zwarts, J.: Forceful prepositions. In: Evans, V., Chilton, P. (eds.) Language, Cognition and Space: The State of the Art and New Directions. Equinox Publishing (to appear)
40. Zwarts, J., Winter, Y.: Vector Space Semantics: A model-theoretic analysis of locative prepositions. Journal of Logic, Language and Information 9, 169–211 (2000)